CHAPTER

5

HARD DISK INTERFACES AND RAID

After reading this chapter and completing the exercises, you will be able to:

- Identify basic physical hard disk components
- Compare physical and logical drives and describe their functionality
- Identify major file systems
- Identify characteristics of the IDE interface and configure IDE cabling and connectors
- Identify characteristics of the SCSI interface and configure SCSI cabling and connectors
- Become familiar with Fibre Channel technology and storage area networking
- Identify and configure various types of RAID

The physical hard disk is the focal point of the enterprise. The operating system, applications, and data are stored on the hard disk, and without hard disk storage, there is nothing for anyone to "do" on the network. Therefore, it makes sense that server storage be protected against failure and that administrators configure storage so that it is as fast as reasonably possible. This chapter opens with a brief inventory of hard disk physical components, differentiates physical and logical drives, and describes the major file systems commonly in use on servers. To operate the hard disk, an IDE or SCSI controller and attachment interface are necessary. Also, you will need to be aware of rules that govern appropriate connections and compatibilities for a given hard disk configuration, especially for SCSI.

Although desktop users commonly use only a single physical hard disk, server administrators often manage dozens of hard disks—most of which require high performance, redundancy, or both using various RAID implementations. When administrators add or replace a drive, it is vital to do so with minimal disruption (if any) to users, and this chapter shows you various ways to do this. Part of

administering server storage is proper maintenance, and you will learn about several such utilities. Finally, you will learn about network storage, including a new technology known as Fibre Channel that provides outstanding performance.

HARD DISK COMPONENTS

Although we use the term "hard drive" or "hard disk," a hard disk is really multiple physical platters inside a sealed, dust-free housing. The manufacturing process is extremely controlled—there is no tolerance for contaminants (such as dust particles)—and manufacturing specifications are extremely tight. Most hard disks include the following major physical components (use Figure 5-1 for reference as you read the rest of this section):

- Disk platters
- Drive heads
- Actuator mechanism

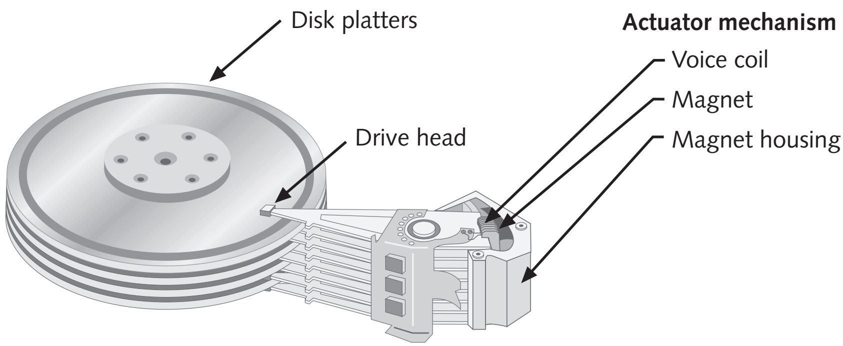

Figure 5-1 Hard disk components

The hard disk also consists of the following components not visible in Figure 5-1:

- Spindle motor
- Logic board
- Connectors
- Jumpers

Though it would be impractical to exhaustively describe each hard disk component, you should be aware of their basic functionality. At the end of this chapter, you will perform an exercise (Hands-on Project 5-1) in which you will visually identify physical hard disk components.

Disk Platters

The **disk platter** is a rigid disk inside the sealed hard disk enclosure. In the past, disk platters have been composed of a metallic aluminum/magnesium alloy, a lightweight and rigid material. On the surface, a syrup containing iron-oxide particles was evenly spread across the disk using centrifugal force. It is this material that stores the magnetic data on the platter. With the cover off the sealed enclosure, you can see that the platter is a brownish or amber color. While you might still have disks in use that were manufactured using this process, it is not implemented in current manufacturing because the oxide medium is soft, making it susceptible to damage if the drive head touches the surface during operation, usually due to a jolt or vibration. (This is known as a **head crash**.) A head crash often (but not always) corrupts the data or compromises the integrity of the recording media.

Current manufacturing procedures have abandoned the metal platter in favor of a glass platter (actually a glass-ceramic composite) because it does not flex as metal can, allowing the platter to be about one-half the thickness of a metallic platter. Thinner platters allow room for more platters in the same drive housing, hence higher-capacity drives. Most current hard disks utilize four platters, but high-capacity drives can use upward of 11. Also, the glass platter is more thermally stable than metal, minimizing the expansion or contraction that occurs with temperature changes.

One way to apply the medium to the platter is to process the platter through a series of chemical baths that leave several layers of metallic film in an electroplating method similar to that used to affix chrome to car bumpers. A better, newer method is to apply a **thin-film** magnetic medium over the glass platter, providing greater density. The thin-film medium is more expensive than electroplating, and it is applied through a process known as "sputtering" in which the material is applied in a continuous, nearly perfect vacuum. Thin-film media are much harder than oxide media. The result is that it is more difficult (nearly impossible) to crash the media. Thin-film media are like a silvery mirror in appearance.

Drive Heads

The **drive head** reads and writes data to the magnetic media on the disk platter. Each platter has two drive heads: one that reads media from the top and another that reads media from the bottom. When stationary, the drive heads are actually parked on the media surface. When the drive is in use and spinning, the air pressure from the movement of the platter separates the drive heads from the surface of the media. Because the drive heads are ganged together in a head rack to a single actuator mechanism (see next paragraph), each head moves across its respective platter in unison with the other heads of the drive.

Note

The air in the drive is specially filtered in two ways. First, a recirculating filter continuously cleans the air as it rotates from platter movement. Because the drive is assembled in a clean room, it does not clean generic dust particles, but small particles from the media scraped from the disk as the heads skid on

the surface during takeoff and landing (spinning up and powering down). Another filter cleans air that enters the drive as a result of barometric equalization with outside air, which is necessary to float the drive heads properly.

The **actuator mechanism** is the mechanical component that physically positions the drive heads at the appropriate location on the disk platter to read or write data. Most actuator mechanisms today use the **voice coil** construction, which derives its name from audio speaker technology using an electrically charged coil. Fluctuations in the electrical charge move the coil to various positions over the platter similar to the way a speaker coil moves to create audible vibrations in a speaker cone. The actuator mechanism has no intelligence of its own in determining the appropriate location; it depends upon the **servo mechanism**, which detects precise cylinder locations on the platter using **gray code**, a special binary code written to the drive by the drive manufacturer that identifies physical locations on the drive. You cannot alter or erase this code, even with the FDISK utility or FORMAT command (these are described in the next section).

Note

Older hard disks required a program to manually park the heads when powered off to avoid having the heads skitter across the surface, causing damage. Today, drives automatically park their heads using a spring and magnetic force. When the drive powers off, the magnetic force of the voice coil actuator dissipates, and a spring drags the head rack to a park-and-lock position.

The physical components described above are relatively durable considering their precise nature. For example, specifications for many drives indicate the MTBF is approximately 1,200,000 hours (that's almost 137 years) and can sustain shock up to 300 Gs (the force of gravity times 300). I once taught a seminar that required me to hot-swap disks on stage. Being highly coordinated and dexterous, I dropped the hard disk from waist height onto a concrete floor. We all gasped, but it turned out fine because the drive suffered no damage and the show went on.

Caution

The environment is still extremely important for hard drives, particularly the temperature. We have already discussed adequate ventilation and cooling, but if your office is in a colder environment, also consider condensation. When a drive comes delivered to you from a cold truck in the middle of winter, allow plenty of time for the drive to warm up (acclimate) to room temperature before using the drive to prevent condensation on internal components. If the drive has been in an environment colder than about 50 degrees F (10 degrees C), allow it to sit at room temperature for several hours before opening the package.

PARTITIONS AND LOGICAL DRIVES

Nearly all operating systems use a storage system that begins with basic Microsoft MS-DOS hard disk partitions. During OS installation, many operating systems (NetWare, for example) might modify the partition or create another partition. To further subdivide and organize storage on partitions, you create logical drives. This section briefly discusses partitions and logical drives.

Creating a Partition

The operating system—whether it is as simple as MS-DOS or as sophisticated as NetWare—requires a defined boundary on the hard disk on which to place its files. The purpose of the partition is to provide this boundary. A partition can be a primary partition or an extended partition (see the section on logical drives). A **primary partition** is a bootable partition on which you can install operating system files. The MS-DOS **FDISK** utility is usually used to create the partition. You can create up to four primary partitions using Windows NT/2000 disk management, but this will be unusual unless you intend to install multiple operating systems on the same server (as in a classroom or lab environment). Typically, you create one primary partition and an extended partition, which can, in turn, contain logical drives.

After creating the partition(s), you must reboot before using the **FORMAT** command to create the file system. If you try to format without rebooting, you will get an error message.

> TIP
>
> Even if you don't use Windows 98 in your environment, I recommend you obtain a Windows 98 startup disk because it has several useful utilities, including FDISK and FORMAT. Booting from the Windows 98 startup disk automatically loads drivers that work with most CD-ROMs, avoiding the need for you to create a customized disk with your specific CD-ROM drivers. I'd also add the **SMARTDRV.EXE** utility and load it prior to installing the NOS because its ability to cache file reads in advance significantly decreases installation time.

Run the FDISK utility by typing FDISK at the MS-DOS command prompt. With more recent versions of FDISK (as in Windows 95B or later), notification that the hard disk is larger than 512 MB appears, and a lengthy prompt asks if large disk support is desired over 2 GB. Nearly always, you want to respond Yes to this option. The series of FDISK menu options is easy to navigate, so I won't detail how to use each one. However, Table 5-1 lists the common options.

Table 5-1 Common FDISK Options

Option	Notes
Create DOS partition or Logical DOS Drive	DOS partition can be a primary or extended partition. Logical DOS drive first requires an extended partition.
Create Extended DOS Partition	Contains logical drives
Set active partition	The BIOS searches for a bootable, active partition, so you must set this to start the operating system. Some operating systems automatically set the partition active at installation.
Delete partition or Logical DOS Drive	Deletes the partition—be careful, you cannot recover data on a deleted partition!
Display partition information	Shows partitions on the drive. Note that any non-FAT file system (described later in this chapter) might appear as a non-DOS partition.

TIP

There are third-party alternatives to FDISK and FORMAT. One alternative is GDISK from Symantec, a command line utility that works much faster and offers more flexibility than FDISK and FORMAT. GDISK accompanies Symantec Ghost, a disk duplication software product. Another alternative is PartitionMagic from PowerQuest, which has a more intuitive graphical interface and is a favorite among administrators for quickly resizing and moving partitions as well as converting them from one file system to another.

Logical Drives

A **logical drive** is a section on the hard disk that appears to the operating system as if it were a separate, distinct hard disk and has its own drive letter. A logical drive requires an extended partition, which can take the place of one of the four primary partitions. You can have a maximum of either four primary partitions or three primary partitions and one extended partition. The sole purpose of the **extended partition** is to store logical drive(s). Logical drives can be lettered up to Z. Within available disk space limitations, you can create as many logical drives in the extended partition as you want. Use FDISK to create the logical drives.

Once the partitions and/or logical drives are created, you must reboot the system and then format the drives. If you do not reboot the system prior to formatting, you are likely to receive an error message. Because the partition is only a storage boundary, you also need to format the partition or logical drive before you can store files on it. It is much like a parcel of land with room for a parking lot. A pavement company can apply asphalt to the entire area available for parking (partitioning), but before people can park, the asphalt must be "formatted" with lines.

Note

Many operating systems include utilities that allow you to create and manage additional logical or primary partitions. For example, Windows 2000 Server includes the Disk Management console with which you can create additional partitions and software RAID configurations (described later in this chapter).

FILE SYSTEMS

A file system is a structure that an operating system uses to name, store, and organize files on a disk. As you administer various operating systems, it is important to understand the basics of each file system. (The following are general descriptions, and are not intended to be exhaustive.)

FAT/FAT32

The Microsoft-based **File Allocation Table (FAT)** file system is compatible with nearly any operating system and uses an invisible table near the top of the partition that points to the physical location of data on the disk. It is the simplest file system, and the network operating systems discussed in this book can all be installed with a FAT file system, though you will often choose to convert to another file system during or after the installation, primarily because of the following FAT limitations:

- *Small volume size*: FAT only supports volumes up to 2 GB in size—tiny even by current home user standards.
- *Large cluster size*: The file system stores data on the drive in 32 KB "chunks" known as clusters or allocation units. If you save a file that is only 2 KB in size, it must use all 32 KB of the cluster, wasting the remaining 30 KB. The space of wasted kilobytes is known as **slack**. Even with hard disk storage at a relatively inexpensive level in recent years, the excessive slack of the FAT file system is undesirable and quickly adds up to several megabytes of wasted space.
- *Limited file size*: The maximum file size is 2 GB. While this seems large, it is woefully insufficient for most corporate databases.
- *Security*: FAT offers no local security. Therefore, any passerby with a boot floppy can fully access the files on the local hard disk. Of course, physical security measures should prevent local access; nevertheless, administrators usually prohibit a strictly FAT file system on the server.

FAT32 is the next (and last) generation of the FAT file system. It overcomes the first two weaknesses of the FAT file system by offering large disk support up to a theoretical 2 TB and using only 4 KB cluster sizes for a significant reduction of slack. However, Windows 2000 only allows up to 32 GB. Also, the maximum file size is 4 GB, but that's still too small for many corporate requirements. Administrators avoid FAT32 because it does not offer security. Microsoft offers a secure file system with the NTFS file system.

NTFS

The **NT File System (NTFS)** is a reference to the Microsoft Windows NT file system. NTFS is compatible only with Windows NT 3.1 or better (including Windows 2000) and is not directly compatible with other operating systems. NTFS volumes offer the following benefits:

- *Large volume size*: NTFS can support extremely large volumes, though the practical limit according to Microsoft is 2 TB.
- *Small cluster size*: NTFS formats clusters at 4 KB each by default on partitions larger than 2 GB, though you can select a cluster size from 512 bytes to 64 KB.
- *Large file size*: File size is limited only by the available drive space.
- *Security*: Unless using hacker's tools, a passerby cannot boot and access files on a local NTFS volume unless they also have a user account that is authorized to access those files. Also, administrators can apply very specific levels of file and folder security that are unavailable on a FAT file system. For example, you could allow a user to save a file or folder but not delete files or folders.
- *Compression*: Though FAT offers compression through Microsoft utilities, these utilities are not considered reliable enough to be practical for server use. NTFS allows you to compress files and folders using a highly reliable compression scheme, conserving disk space.
- *Data integrity*: NTFS includes mechanisms designed to ensure that data is properly and completely written to the drive.
- *Windows 2000 NTFS features*: Under Windows 2000, a variety of new features appear in NTFS 5. You can encrypt files and folders with a nearly unbreakable encryption scheme, and set quotas so that users do not abuse file storage privileges. Besides the file system itself, the operating system offers many useful capabilities for managing partitions and files, including the **Distributed File System (Dfs)** to deploy what appears to be a single directory structure over multiple physical file servers. Windows 2000 also has offline storage to migrate seldomly used files to a slower and less expensive storage medium such as tape.

HPFS

IBM uses the **High Performance File System (HPFS)**, which bears many similarities to NTFS in basic structure. In fact, Windows NT was originally supposed to be named OS/2 Version 3.0 in 1993. However, IBM is not currently pursuing further developments to HPFS, and so NTFS has surpassed HPFS in many respects. HPFS is designed to provide local security similar to NTFS, but also coexists with an MS-DOS partition, allowing you to dual-boot between either MS-DOS or OS/2. OS/2 includes the **Journaled File System (JFS)**, which contains its own backup and recovery capability.

Using an indexing system and log to corroborate file changes, JFS can interoperate with the operating system to repair corrupt files. HPFS supports a maximum partition size of 2 TB and 2 GB file sizes, which was a staggering size when OS/2 was at its peak. HPFS also offered long file name support (up to 254 characters), 512-byte clusters, intelligent link tracking (you can move a file and the "shortcut" automatically points to the new location), and efficient file distribution on the hard drive to minimize fragmentation. Along with OS/2, Windows NT 3.x can also access HPFS partitions.

Note

Though OS/2 is much less popular than other NOSs, the CompTIA Server+ exam requires knowledge of its basics.

Linux/UNIX File Systems

There are dozens of different "flavors" (versions) of UNIX in use. For the most part, this book and the CompTIA Server+ exam gravitate toward Linux as the specific implementation of UNIX. UNIX implementations generally use the **UNIX File System (UFS)**, the **Network File System (NFS)**, or **AFS**, which stems from Carnegie-Mellon's Andrew File System. Linux often uses the **Filesystem Hierarchy Standard (FHS)**, which is more a directory structure than a file system (see *www.pathname.com/fhs* for the complete standard). Finally, you can also create Linux-specific partitions—Linux swap, Linux native, and Linux RAID, each of which is addressed further in Chapter 8.

UNIX file systems have a higher administrative learning curve—you must use an arcane command line interface. However, several interfaces (such as X Windows) now allow you to perform many UNIX administrative functions using a GUI. Also, various flavors of Linux characteristically include a GUI.

NetWare

The traditional **NetWare file system** competes directly with Microsoft's NTFS, and therefore offers similar features: You can use very large volumes and files, and the cluster size is efficient. Using the traditional NetWare file system, you can create a NetWare **volume**, which is a collection of files, directories, subdirectories, and even partitions. You can combine separate partitions that together comprise a volume. File storage is also efficient because NetWare volumes can use suballocation. A NetWare volume can subdivide a block (a Microsoft cluster is a NetWare "block") to minimize slack space. Unfilled blocks are subdivided into 512-byte suballocation blocks, which can then be used to store data from one or more other files.

Like NTFS, NetWare offers file compression. You can manually activate compression by flagging files and directories with the IC (immediate compress) command or by using the SET command at the server console to configure a file inactivity delay, after which compression occurs automatically. NetWare offers a distributed file system and offline file storage similar to Windows 2000 Dfs and offline file storage, respectively.

Note

NetWare 5.x uses two file systems: the "traditional" NetWare file system described above and a new NSS file system described later in this section. Be careful not to use an "NFS" acronym for the traditional NetWare file system because it is easily confused with the NFS of the UNIX Network File System.

NetWare 5.x also includes additional features in its optional **Novell Storage Service (NSS)** file system. The purpose of NSS is to increase performance and total storage capacity. NSS offers several improvements over the traditional NetWare file system:

- *Large files*: Instead of the traditional 2 GB limitation, files can now be up to 8 TB each! Also, you can store *trillions* of files in a single directory, compared to the 16 million entries per volume under traditional limitations.
- *Performance*: Large files typically take a long time to open, but NSS provides rapid access regardless of size. Also, mounting a volume (preparing it for use) is much faster.
- *Flexible storage management*: You can create up to eight NetWare partitions on a single disk and create unlimited volumes per partition.

Note

You can't just install an NSS file system; you have to first have the traditional NetWare file system and then add NSS.

THE IDE INTERFACE

Integrated Drive Electronics (IDE) can refer to any hard disk with an integrated controller. However, it is more technically accurate to apply the term **ATA (AT Attachment)** to what we usually call an IDE drive, because the drive plugs into the 16-bit ISA bus known as the AT bus. In spite of this technicality, this book and most references use the terms IDE and ATA synonymously. Recall that an IDE drive includes the controller in the circuitry attached to the drive. When people colloquially say that they are installing an "IDE controller," what they really mean is that they are installing a **host adapter**, the more accurate term for what is usually referred to as an IDE or SCSI hard disk controller. The host adapter is the physical interface between the hard disk and the computer bus. As you can see, disk interface terminologies are replete with inaccuracies (and we have only gotten started). As we delve into the IDE interface and its variations, you'll learn synonyms for various interfaces.

Note

Although the ATA interface attaches to the 16-bit ISA bus, it is not a performance concern because even two of the fastest IDE drives cannot saturate the ISA bus.

Other devices besides hard disks can plug into the ATA interface, usually CD-ROM drives and tape drives. These devices require a variation on the ATA specification known as the **ATAPI (ATA Packet Interface)** specification. The SCSI interface, discussed later in this chapter, can also accept other devices. For this chapter's purposes, however, we are only concerned with hard disks for both interfaces.

> Note: Many times, an inaccuracy in terminology develops when a hard disk drive manufacturer coins a term to help sell its product. For example, ATA-2 is often referred to by its marketing term EIDE.

The next few sections will describe variants of the ATA interface. Early generations of the ATA interface will be discussed in less detail, because you are not likely to find them in servers. However, some information about them is relevant because standards introduced with them carry through to current ATA specifications.

ATA-1

The following major features characterize the **ATA-1** standard:

- Signal timing for DMA and Programmed I/O (PIO), which utilizes the processor to handle disk transfers but is superceded by DMA and Ultra-DMA
- 40/44-pin cable connections (44-pin connections use four more pins to supply power to notebook hard drives)
- Determination of master, slave, or cable select using jumpers (discussed later in this chapter)
- Transfer rate of 3.3 MBps to 8.3 MBps depending on the PIO or DMA mode

ATA-2

The following major features characterize the **ATA-2** standard:

- Large drive support for up to 137.4 GB (previously 8.4 GB)
- Faster PIO and DMA transfer specifications
- Power management support
- Removable device support
- PCMCIA (PC card) support
- Reports drive characteristics to software (useful for Plug and Play)
- Transfer rate of 8.3 MBps to 16.6 MBps depending upon the PIO or DMA mode

Before proceeding, let's sort out some more terms. ATA-2 is synonymous with the unofficial marketing terms Fast-ATA, Fast-ATA-2, and Enhanced IDE (EIDE). EIDE has become one of the most accepted terms, and it applies in a general way to ATA-2 or better. The true published specification is AT Attachment Interface with Extensions, but people seldom use that title.

In order to support drives larger than 8.4 GB, you must also have an operating system that is capable of recognizing larger drives. The following operating systems support drives larger than 8.4 GB:

- *Windows 95B or later*—Using FAT32, you can format drives up to 2 TB.
- *Windows NT 4.0*—This OS supports larger drives out of the box, but not on the bootable drive. For the bootable drive, you must first apply Service Pack 4 to enable large drive support.
- *Windows 2000*—This OS provides native support for larger drives.
- *OS/2 Warp*—With a Device Driver Pack upgrade, a boot partition can be only as large as 8.4 GB. If you use HPFS, OS/2 supports up to 64 GB.
- *NetWare 5.x*—This OS provides native support for larger drives.

ATA-3

The following major features characterize the **ATA-3** standard:

- Includes Self Monitoring and Reporting Technology (S.M.A.R.T.), a predictive technology that enables the operating system to warn of a device's degradation. S.M.A.R.T. has its basis in preceding technologies known as Predictive Failure Analysis (IBM) and IntelliSafe (Compaq). Drives might use one of the three technologies, or a combination. For example, some Compaq systems use both IntelliSafe and S.M.A.R.T. If the hard disk begins to show signs of failure, you will see messages in the server's system log, RAID log, or in a vendor-supplied monitoring and reporting utility.
- Optional security mode that protects access to the drive with a password
- Transfer rate of 11.1 MBps to 16.6 MBps depending on the PIO or DMA mode

Note

S.M.A.R.T. has made its way into the SCSI world of hard disks as well (see more about SCSI later in this chapter). This significantly adds to the administrator's ability to monitor the health of internal and external SCSI disks and RAID configurations.

ATA-4

The following major features characterize the **ATA-4** standard:

- Addition of the ATAPI standard to attach other types of devices
- Advanced power management

- Specification of an optional 80-conductor, 40-pin cable-select cable to reduce noise
- Improved BIOS support for a theoretical capability of 9.4 trillion gigabytes, though the actual ATA standard is still limited to 137.4 GB
- Ultra-DMA (UDMA) support, increasing the transfer rate to 33 MBps

ATA-4 is probably the most groundbreaking standard in terms of current IDE performance. It introduced a level of performance (which was increased even further with ATA-5 and ATA-6) that had been formerly available only on the SCSI interface. Other terms for ATA-4 include Ultra-DMA and Ultra-ATA. In reference to the transfer rate, you might also see UDMA/33 or Ultra-ATA/33. In previous ATA implementations, data is transferred once each clock cycle. Ultra-ATA differs in that data transfers twice for each clock cycle, once at the rising edge and once at the trailing edge. Ultra-DMA also adds a cyclical redundancy check (CRC) to ensure the integrity of data.

To support Ultra-ATA/33 and later, a compatible drive, BIOS, operating system, and host adapter interface must be in use. In the BIOS, most manufacturers now include an artificial "32-bit" transfer. Recall that IDE operates at 16 bits on the ISA interface, even though the host adapter is usually a PCI card. The BIOS now includes functionality that allows for two 16-bit transfers to occur at once, hence the "32-bit" transfer.

Note

Under ATA-4 and higher, a single drive on the IDE cable must be at the end of the cable (no "stub" allowed). Otherwise, signaling problems can occur. Under earlier ATA versions, a stub was OK but inadvisable.

TIP

If you have any older ATA-1 through ATA-3 drives, don't throw them away just because current standards are ATA-4 or better. The ATA specification requires successive ATA iterations to be backward compatible. For example, you could still attach an ATA-1 drive to an ATA-4 or higher host adapter for cheap (but slower) storage. Note that as a general rule, if you mix ATA standards on the same cable, both devices operate at the performance level of the slower standard.

ATA-5

The following major features characterize the **ATA-5** standard:

- 80-conductor cable required (as opposed to optional) in order to achieve the maximum transfer rate. You can use a standard 40-pin cable but only at a maximum transfer rate of 33 MBps.
- Added to the ATA-5 specification is an IEEE-1394 (FireWire) link that allows use of an ATA drive on the FireWire interface. (**FireWire** is an extremely fast bus allowing up to 63 connected devices and up to 3200 Mbps throughput in the latest version.)

- Transfer rate of 66 MBps, achieved by reducing setup times and increasing the clock rate. Later implementations of ATA-5 achieve 100 MBps under the marketing title of Ultra-DMA/100. This came about as a result of manufacturers that could match the 100 MBps transfer rate of the ATA-6 standard but did not want to wait for the completed standard.

Starting with ATA-5 and ATA-6 (see next section), you are most likely to have to add an adapter to achieve the maximum throughput because it takes several months and sometimes over a year for motherboard IDE interfaces to catch up to the latest ATA standard.

ATA-6

At this writing, the **ATA-6** standard was still under development. The current draft of this standard is the 475-page 1410D Revision 1b (*www.t13.org/project/d1410r1b.pdf*). However, this standard should include all features of previous ATA standards plus a formalized 100 MBps specification.

Table 5-2 summarizes the IDE standards covered in the previous sections.

Table 5-2 IDE Standards

IDE Standard	Also Known As	Performance
ATA-1		3.3–8.3 MBps
ATA-2	AT Attachment Interface with Extensions, Fast-ATA, Fast-ATA-2, and Enhanced IDE (EIDE)	8.3–16.6 MBps
ATA-3	EIDE	11.1–16.6 Mbps
ATA-4	Ultra-DMA, Ultra-ATA, UDMA/33, or Ultra-ATA/33	33 MBps
ATA-5	Ultra-DMA/100 for 100 MBps implementations	66–100 MBps
ATA-6		100 MBps

ATA Cable

The standard ATA cable connecting the drive to the host adapter is a 40-pin ribbon cable (see Figure 5-2). To prevent incorrect connections, a cable key (protruding notch) on the cable matches a corresponding gap in the IDE connection on the hard drive. The cable should have one striped wire (usually red, sometimes blue) that indicates Pin 1. Pin 20 is not used, and is usually absent from the drive, and a corresponding block in position 20 appears on the cable, also preventing backward insertion (see Figures 5-2 and 5-3). Orient the connection so that Pin 1 on the cable is adjacent to the power cable connection. Because the cable is not shielded, you are limited to a length of 18 inches (457.2 mm). A longer cable could be more sensitive to timing and electrical noise issues, resulting in data corruption. The host adapter connected to the other end should also have a cable key, stripe, and absent Pin 20, and most manufacturers include a marking on the PCB (printed circuit board) that indicates Pin 1.

The cable length of 18 inches (457.2 mm) is usually plenty for a desktop workstation, but in full-size server tower cases in which the controller is farther away from the drive bays, this might cause a problem. If the ATA host adapter is a card (as opposed to integrated on the motherboard), move it as close as possible to the drive locations. Otherwise, if you need longer lengths, you can obtain a longer, custom-made IDE cable, but at the risk of performance degradation and data loss.

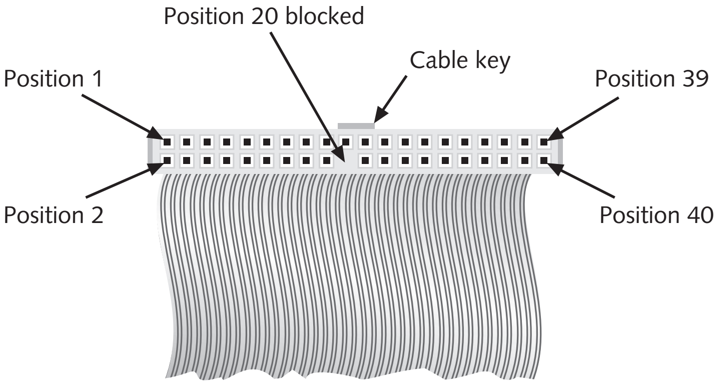

Figure 5-2 Standard 40-pin IDE connector

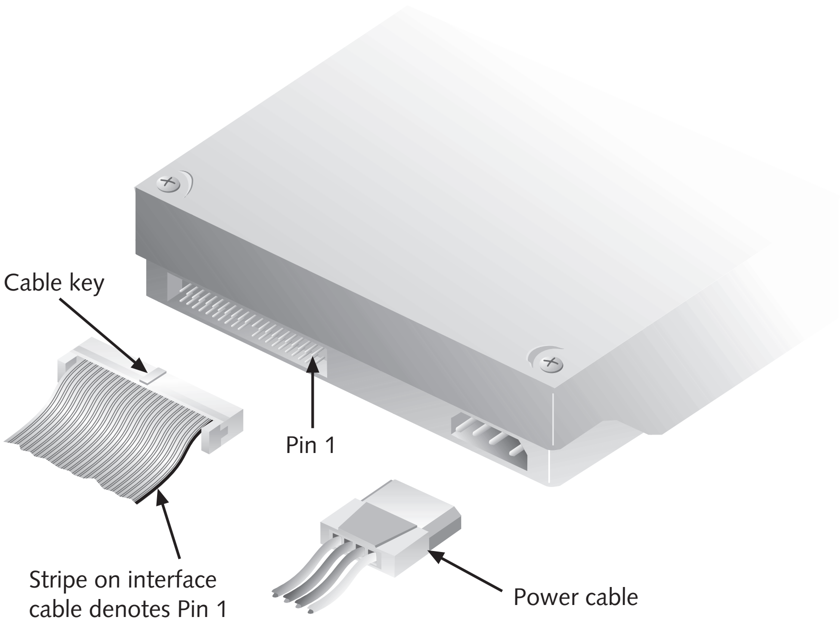

Figure 5-3 Correct orientation of the IDE cable on the drive

You can connect a maximum of two devices to a single ATA connector. The cable and its dimensions appear in Figure 5-4. Also, 80-conductor cables are identical except for an additional 40 grounded conductors. At one end of the cable, an IDE port connector is blue. Two more connectors appear on the cable—the first one is at the opposite end and is black and the middle one is gray. Connect the first hard disk to the black end connector. If you only have one hard drive, also make sure you connect it to this one and not the middle connector. Otherwise, the dangling connector on the end (the "stub") is not well terminated. Although it is not against ATA specifications, attaching a single drive to the middle connector is not recommended. Connect a second hard drive to the middle gray connector.

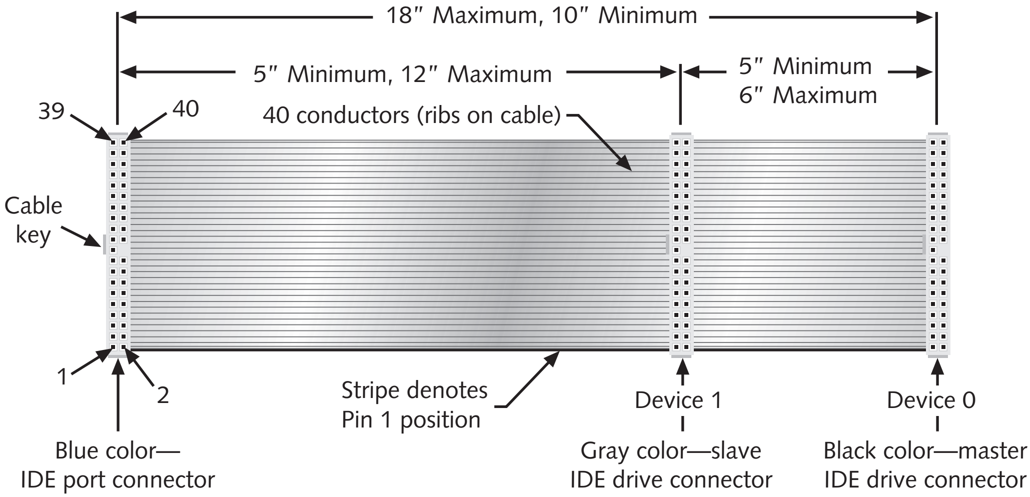

Figure 5-4 ATA 40-pin color-coded cable

> **TIP** You might find some systems, usually very cheap desktop PCs, which have none of the color-coding that you use to make proper connections. As a result, you have to guess which way to orient the cable. Fortunately, if you attach it in an upside-down orientation, it does not damage the drive or host adapter because the cable is not powered—it only carries data signals. If you plug a non-color-coded cable in and the drive doesn't work, just turn the cable over.

As an option, you can use the 80-conductor 40-pin cable with ATA-4, but it is required with ATA-5 and ATA-6 (see Figure 5-5). An 80-conductor 40-pin cable leads to the question: If there are 80 conductors, why isn't it an 80-pin cable? The additional 40 conductors are connected to ground only, and do not have a corresponding pin on the drive. Remember that an IDE cable is not shielded in any way, and is susceptible to electrical noise, usually crosstalk from adjacent conductors (hence the ground wires between signal wires). The extra grounded conductors absorb much of the electrical noise that would otherwise defeat the added performance of ATA-4 and better drives.

TIP While optional on ATA-4, I recommend the 80-conductor cable for any version of ATA because of better signal quality.

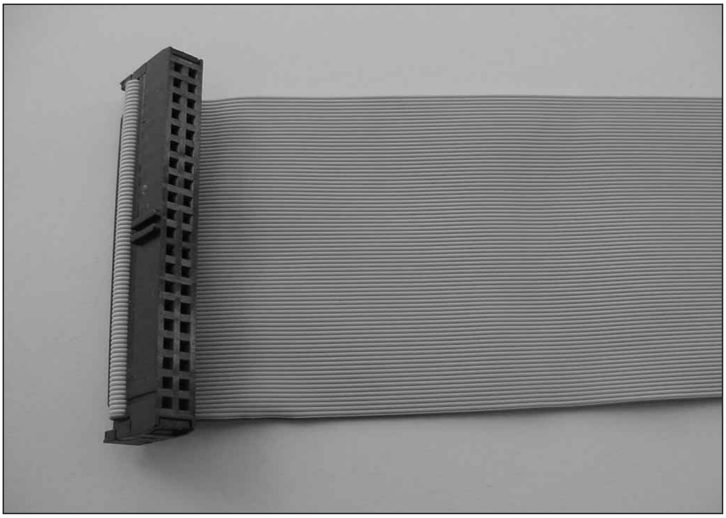

Figure 5-5 An 80-conductor, 40-pin cable

TIP If you connect a faster device such as a hard disk and a slower device such as a CD-ROM to the same IDE cable, the hard disk performance will suffer while waiting for the CD-ROM to finish its tasks. Because IDE is not capable of simultaneous I/O, one device must wait for the other before performing tasks. In this case, it would be better to obtain another ATA host adapter and keep slower devices (such as CD-ROMs and tape drives) on one IDE connector and hard drives on the other.

Master, Slave, and Cable Select

Because you can place two drives on an IDE host adapter, there must be a determination as to which one is the **master** and which one is the **slave**, especially because the master receives a drive letter assignment from the operating system first, and is also the device on which a boot record must be found. Otherwise, the master and slave drives are equivalent despite the implication of the names. You can use two methods to specify which drive is master and which is slave: First, you can set the jumper on the drive to indicate a master or slave setting. The drive manufacturer specifies these jumper settings in the drive manual or online documents. For example, Figure 5-6 shows the settings downloaded from Maxtor for the Diamondmax 80 Ultra-DMA/100. Most manufacturers also include an indication on the PCB or somewhere on the drive case

as to which jumpers to set for master, slave, or cable select. Although you might need to change the jumper setting to specify master or slave for some older IDE drives, newer drives allow **cable select**, which means that the drive's position on the cable indicates whether it is a master or slave. Place the drive on the end of the 80-conductor cable to make it a master, and on the middle connector to make it a slave.

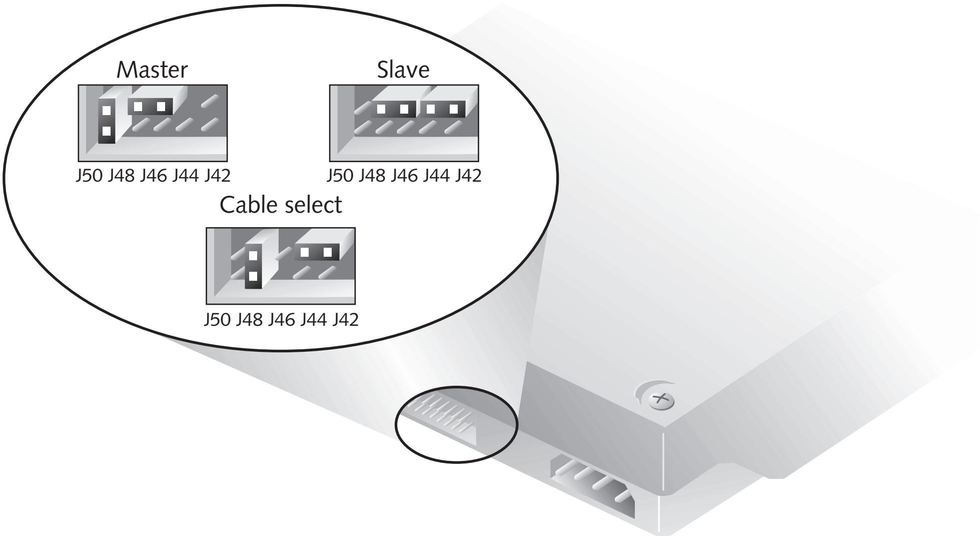

Figure 5-6 Jumper settings for master, slave, or cable select

TIP: Cable select does not work on a standard 40-pin cable unless you use a specially configured cable. To make matters more complicated, the master/slave position on a 40-pin standard cable is reversed—the master drive is on the middle connector and the slave drive is on the end connector.

Note: It's not as much an issue as it used to be, but sometimes if you mix two drives from different vendors on the same IDE channel, you may experience problems. For example, the system might be unable to detect one of the drives, or if the drives are different sizes, report both drives at the smaller size. If you experience problems of this kind and are certain jumpers and locations on the IDE cable are correct, consider placing drives of the same manufacturer (and perhaps the exact same model) on the same cable.

The Pros and Cons of IDE/ATA

Current ATA implementations have some compelling advantages for use in the server world, but also have several disadvantages. If the disadvantages outweigh the advantages, and budgeting is sufficient, you will probably choose the more expensive SCSI interface (discussed later in this chapter).

Pros

- *Inexpensive*—Both the drives and host adapters are very inexpensive; even administrators on the tightest of budgets can probably afford a few new IDE drives.
- *Reasonable performance*—In less demanding situations where the server is not a heavily utilized file server, the reasonably good performance might be acceptable considering the low cost.
- *Simple configuration*—Plug it in and format the drive. That's pretty much it, and you don't usually have configuration complications, which you might experience with SCSI, for example (discussed later in this chapter).

Cons

- *Only two drives*—Two drives might seem like a lot for a workstation, but a server often connects to many more drives for purposes of performance, redundancy, and total available storage.
- *Slower device throttling*—Mixing slower devices with faster ones can drag the performance of the faster one.
- *Cable length*—18 inches (457.2 mm) is very limiting in a full-size server.
- *No simultaneous I/O*—Only one IDE device can operate at a time on a single IDE channel. SCSI, which does permit simultaneous I/O, might be much more attractive in terms of performance.
- *No native redundancy*—You cannot take a standard ATA drive and host adapter and configure hardware RAID. However, software RAID is available on certain operating systems (Windows 2000, for example). Also, some manufacturers such as Promise Technology Inc. (*www.promise.com*) are making ATA adapters that allow you to configure several levels of RAID. (See more about RAID later in this chapter.)

TIP

Although simultaneous I/O is not supported on a single IDE channel, you can place two drives on separate IDE channels. In this configuration, each drive can operate simultaneously.

- Despite advertised transfer rates, most of the time you only get that rate from cached data. Data retrieved fresh from an Ultra-ATA/100 hard disk can usually only be sustained at about 40 MBps.

TIP

More information about ATA/IDE can be found at *www.ata-atapi.com* and especially *www.pcguide.com*.

THE SCSI INTERFACE

SCSI (Small Computer System Interface) technology, like ATA, has seen multiple generations of specifications, and includes an abundance of technical details beyond the scope of this book. However, as a server administrator, you must be aware of the major SCSI standards in order to provide the best available compatibility solution with the best possible performance. Properly configuring your SCSI hard drives also avoids complex troubleshooting.

As with IDE, SCSI drives include the controller circuitry directly on the hard disk assembly (HDA). In fact, some sources claim that hard disk manufacturers nearly mirror their drives in all physical respects and only change the circuitry to distinguish IDE and SCSI, often differentiated by only a single chip. What most people call the SCSI "controller" is actually the SCSI host adapter—similar to the inaccurate reference to an IDE "controller." However, although an IDE device communicates directly on the system bus, a SCSI device first communicates with the CPU through the SCSI host adapter.

At the end of this section, you will see a listing of the various pros and cons of using the SCSI interface. Generally, however, the main "pro" is that it performs exceptionally well in multiple disk configurations. The main "con" is that when you combine the more expensive SCSI hard disks with the more expensive SCSI host adapters, the price tag is significantly higher than an IDE configuration.

SCSI standards range from SCSI-1 to SCSI-3 and encompass several variations in between. You are unlikely to find SCSI-1 in modern server implementations, but you need to know some basics about it according to the stated CompTIA Server+ exam objectives. SCSI-2 (especially the Fast and Wide iterations) and some SCSI-3 variants are more likely to be found in servers today. With most SCSI standards, you will see which cables and/or terminators are required, all of which are explained in more detail later in this chapter.

Note

Each succeeding version of the SCSI standard is backward compatible with all preceding versions. In theory, you could put a SCSI-1 hard disk on a SCSI-3 host adapter. However, such drastic implementations are not recommended because of complicated cabling adapters, termination incompatibilities, and performance limitations.

A more complete description of the various signaling types, cables, and connectors appears after the sections on individual SCSI versions.

SCSI-1

SCSI-1 started off as just "SCSI" but was later renamed SCSI-1 to avoid confusion with successive SCSI standards. The following major features characterize the SCSI-1 standard:

- 8-bit **parallel bus**, meaning that multiple wires on the cable can transmit data at the same time. This is what allows SCSI to use simultaneous I/O.

- 50-pin Centronics-style external connector and low-density pin header internal connector
- Single-ended (SE) transmission
- Passive termination (the simplest type of termination, but also the least reliable; see later section on SCSI termination)
- Optional bus parity checking
- 5 MHz operation
- Transfer rate of 4 MBps (asynchronous) or 5 MBps (synchronous)

SCSI-1 is now considered obsolete and has been retired by ANSI and replaced with SCSI-2.

SCSI-2

SCSI-2 is essentially SCSI-1, plus the following optional features:

- **Fast SCSI** operating at 10 MHz instead of 5 MHz
- **Wide SCSI** utilizing 16-bit transfer instead of 8-bit (which is "narrow")
- 50-pin high-density connectors
- Active termination (see later section on SCSI termination)
- High Voltage Differential (HVD) is used to extend bus length. HVD is now considered obsolete in favor of Low Voltage Differential (LVD). (HVD and LVD are covered later in the chapter.)
- **Command queuing** allows the host adapter to send as many as 256 commands to the drive. The drive stores and sorts the commands for optimum efficiency and performance internally before responding to the host adapter. Multitasking operating systems such as OS/2, Windows NT, and Windows 2000 can take advantage of command queuing.
- Transfer rate of 10 MBps at 16 bits and 5 MHz (Wide SCSI), 10 MBps at 8 bits and 10 MHz (Fast SCSI), or 20 MBps at both 10 MHz and 16 bits (Fast and Wide)

SCSI-3

SCSI-3 is the most confusing of the SCSI standards, because once again, manufacturers have used misleading marketing language. In addition, the SCSI-3 standard is not a complete standard of its own. Instead, it is more a collection of documents covering new commands, electrical interfaces, and protocols. A manufacturer could comply with only one of the major SCSI-3 additions and still label the product "SCSI-3." Nonetheless, subdividing SCSI-3 into several smaller standards helps SCSI-3 implementations to develop more quickly than waiting for the entire standard to be published and approved.

SPI SCSI-3 Parallel Interface

SPI is the SCSI-3 Parallel Interface, better known by the marketing terms Ultra SCSI or Wide Ultra SCSI. Along with this standard is the separate **SCSI Interlock Protocol (SIP)** defining the parallel command set. This standard was published as three documents offering the following features:

- 10 MHz bus speed, which is really no faster than SCSI-2. This is an example of how a manufacturer could use the 10 MHz bus speed (Fast SCSI) and call it SCSI-3 even though it offered no performance advantage.
- Fast-20 offers transfer rates up to 40 MBps using 20 MHz signaling.
- 68-pin P-cable and connectors for Wide SCSI

SPI-2

Also known as Ultra2 SCSI and Wide Ultra2 SCSI, **SPI-2** is characterized as follows:

- Single Connector Attachment (SCA-2) connectors. A successor to the problematic original SCA connector, this is the connection type that is mostly used in a chassis that contains several hot-swappable SCSI drives.
- Fast-40 40 MBps transfer rate on a narrow (8-bit) channel or 80 MBps on a wide (16-bit) channel. LVD is required for these data rates.
- LVD signaling to replace the previous SE signaling is required to achieve the faster throughput of Ultra2 or Ultra2/Wide speeds (40/80 MBps). You can use an SE device on the LVD SCSI chain, but doing so switches the chain to SE mode, which throttles the performance of the chain to SE-level performance at a maximum of 40 MBps and shortens the total cable length to as little as 5 feet (1.5 meters) in Fast-20 mode.
- 68-pin Very High Density Connector (VHDC) makes the connectors and ribbon cables smaller. This is important in SCSI because the older connectors and ribbon cables could be quite large and cumbersome.

SPI-3

SPI-3 was still in the draft stage at the time of this writing, although SCSI manufacturers have already implemented many of its features, including:

- Improving the earlier parity check, the **cyclical redundancy check (CRC)** is a calculation used by the sending device based on the data in the packet. The data arrives at the destination target and another calculation is performed using the same "formula." If the calculation in the packet matches the calculation performed by the destination device, the data is complete and considered error free.
- Domain validation improves the robustness of data transfer. In the past, the host adapter would send an inquiry to the SCSI device to determine its supported

transfer rate. If the interconnection between host and device did not support the full transfer rate, then the device became inaccessible. With **domain validation**, the determined transfer rate is tested, and if errors occur, the rate is incrementally reduced and again tested until no errors occur.

- **Double transition (DT) clocking** transmits data on both the rising and falling edges of the clock. On a 16-bit, 40 MHz bus, this yields a transfer rate of 160 MBps.
- **Packetization** reduces the overall communication method to transfer data. Previously, data was transferred over the SCSI bus using a series of phases to set up and transfer data. Packetization streamlines this process by combining the process into a packet, reducing overhead.
- **Quick Arbitration and Selection (QAS)** eliminates the previously required **arbitration** method in which devices contend for control of the bus. Much of the prioritization is based on the device's priority level based on its **SCSI ID**, a unique number for each SCSI device. Because no data transfer can occur during the time that arbitration takes place, it adds overhead to the bus. QAS reduces overhead by reducing the number of times that arbitration must occur and by allowing a device waiting for bus access to do so more quickly.

SPI-3 is also known as Ultra3 SCSI.

Ultra160 and Ultra160+

The five features listed for SPI-3 sometimes cause a problem in the way a SCSI product is marketed. A manufacturer might include only one of the five items and call its product SCSI-3. Most manufacturers prefer to have more stringent requirements, so for the Ultra160 and Ultra160+ standards, the product must include:

- DT clocking at 160 MBps
- Domain validation
- CRC

QAS and packetization are optional for Ultra160, but required for Ultra160+.

SPI-4

SPI-4 was still in draft form at the time of this writing, but most hard disk and host adapter manufacturers have products using the standard's 320 MBps data rate. This data rate is accomplished by doubling the bus speed from 40 MHz to 80 MHz and using DT clocking. Manufacturers are calling this standard Ultra320. Other special features are not yet confirmed.

5

Note: For more information on this emerging standard as well as other SCSI standards, see *www.t10.org*.

Table 5-3 summarizes SCSI standards and performance.

Table 5-3 SCSI Standards

SCSI Standard	Also Known As	Performance
SCSI-1 Async	Asynchronous	5 Mhz/8-bit/4 MBps
SCSI-1 Fast-5	Synchronous	5 MHz/8-bit/5 MBps
SCSI-2 Fast-5/Wide	Wide	5 MHz/16-bit/10 MBps
SCSI-2 Fast-10	Fast	10 MHz/8-bit/10 MBps
SCSI-2 Fast-10/Wide	Fast/Wide	10 MHz/16-bit/20 MBps
SPI (SCSI-3) Fast-20	Ultra	20 MHz/8-bit/20 MBps
SPI (SCSI-3) Fast-20/Wide	Ultra/Wide	20 MHz/16-bit/40 MBps
SPI-2 (SCSI-3) Fast-40	Ultra2	40 MHz/8-bit/40 MBps
SPI-2 (SCSI-3) Fast-40/Wide	Ultra2/Wide	40 MHz/16-bit/80 MBps
SPI-3 (SCSI-3) Fast-80DT	Ultra3 (Ultra160)	40 MHz/16-bit/160 MBps
SPI-4 (SCSI-3) Fast-160DT	Ultra320	80 MHz/16-bit/320 MBps

SCSI Configuration

Configuring SCSI devices requires knowledge of the various types of SCSI, cables, capabilities, and terminators. One of the most common sources of frustration and troubleshooting is an improperly configured SCSI bus. When configuring SCSI from scratch, first consider the budget available to you. SCSI can quickly become cost-prohibitive, especially with several high-performance hard drives. Closely associated with the cost is the level of performance and amount of storage you want. Of course, the more you want, the more you pay.

TIP: Administrators tend to gravitate toward SCSI drives for most purposes, but if you are using only a single disk in your system, it is much more cost-effective to use ATA-4 or better because the performance is reasonably close to SCSI and the cost is a comparative bargain.

SCSI Cables and Connectors

Although you won't see every possible combination of cables and connectors listed here, you should be aware of those that are most commonly used. A narrow bus width (8 bits) is physically and logically smaller than the wider 16-bit bus. The 8-bit, 50-conductor

cable is also known as **"A" cable** and the 16-bit, 68-conductor cable as **"P" cable**. (Other cables have been proposed, but have not come to fruition because the standards never caught on with manufacturers.)

The "A" cable is a low-density cable with a Centronics-style connector (see Figure 5-7) and 50 conductors. The "A" cable is used only in 8-bit (narrow) implementations. If you have both narrow and wide cables, you can continue to use both standards with special adapters such as the one in Figure 5-8, which reduces a 68-conductor cable to a 50-pin connector.

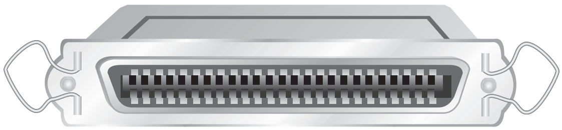

Figure 5-7 50-pin, low-density "A" cable connector

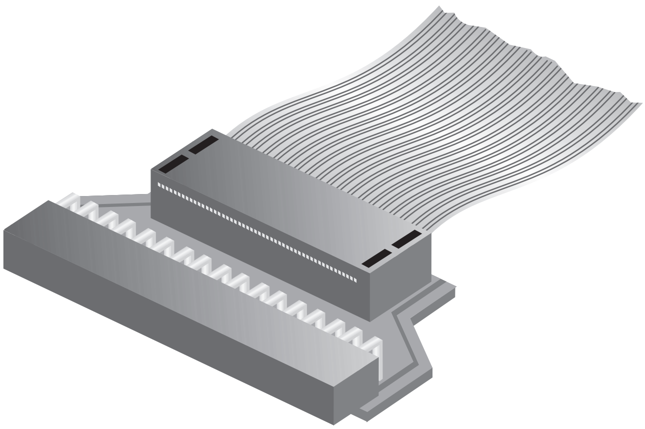

Figure 5-8 50-pin connector on a 68-conductor cable

Note

All this talk about 8-bit and 16-bit widths leads to the question: What about 32-bit widths? While it sounds like a good idea (and probably is), manufacturers have scoffed at a 32-bit cable standard because of exorbitant development and manufacturing costs.

Thankfully, in the SCSI-2 standard, a high-density 50-pin SCSI connector was introduced to save space, because large SCSI connectors such as the Centronics type are cumbersome and can quickly take up space in a cabinet. Also, instead of using the wire latches on either side of the connector to attach and remove the connector, you use squeeze-to-release clip locks.

On the SCA-2 connector, the chassis has the female connector and the drive has the male connector. (Recall that this is the connector that allows hot-pluggable hard drives.) On the outside edges of both connectors are advanced grounding contacts that allow you to

pull out or plug in a SCSI drive without negative electrical consequence (see Figure 5-9). Because you can't see the actual connection take place inside the chassis, this is known as **blind connector mating**. The connector provides both signal and power, and is an 80-pin connector offering only wide (not narrow) SCSI.

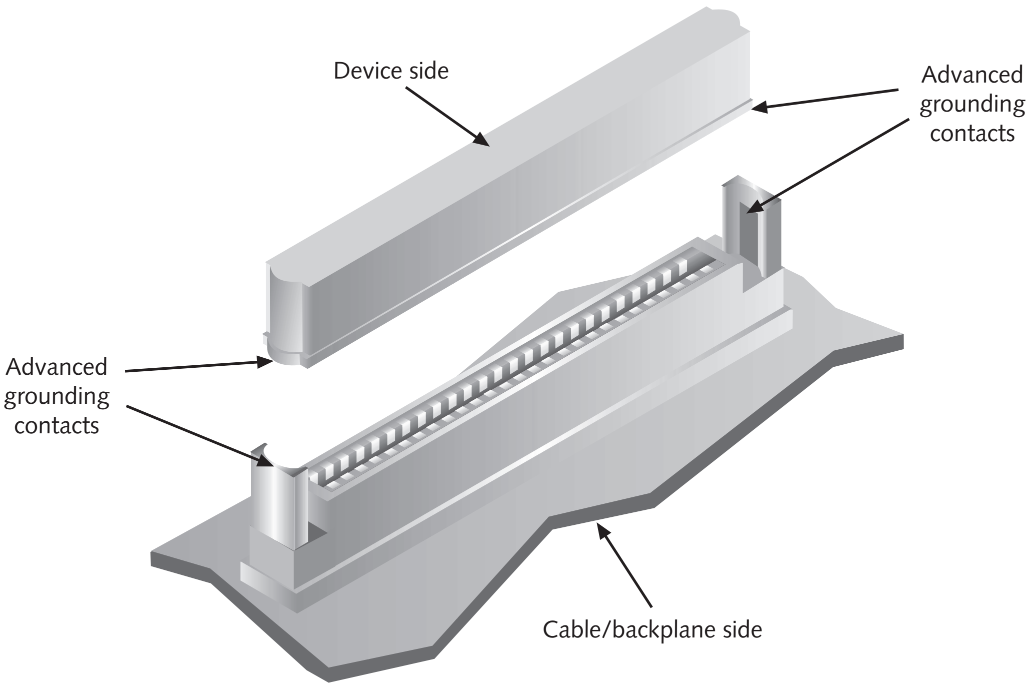

Figure 5-9 SCA-2 SCSI attachment

With SCSI-3, you most often find the wide SCSI cable using a 68-pin connector both internally (inside the case) and externally (outside the case and connected to the server's SCSI host adapter).

Some SCSI connectors, though rare, are 25-pin connectors created for economy of both design and dollars. This connector is streamlined compared to the 50-pin cable, and accomplishes this design mostly by removing grounding—an inadvisable practice. Recall that grounding is a contributor to signal integrity. Because the signal quality is not as good as standard 50-pin or greater SCSI cable, and because the connector is visually identical to the commonly used DB25 parallel cable connector, this style of cable/connector is not recommended. Accidentally plugging printers into the SCSI connector or vice versa can damage equipment.

Connectors are designed for either internal or external use. External cables are usually round with thumbscrews or clip-lock connections on the connector. The round cable is highly shielded and well engineered against signal degradation and interference. As a result, it is also relatively expensive. Internally, you will see ribbon cable in any of several types.

Signaling

In order to obtain the correct accessories for your SCSI chain, you must first know what kind of signaling is in use on the bus. **Signaling** generally describes transmission of data using electrical impulses or variations. These electrical transmissions represent data that the sender originates and the receiver translates based upon a mutually agreed-upon method. There are three types of signaling on the SCSI bus, detailed in the following sections. You must be careful about mixing devices intended for one type of signaling with devices intended for another type of signaling, because one or more devices might stop functioning or become damaged. Also, you must know the signaling before you can properly terminate a SCSI bus.

SE

Single-ended (SE) signaling, the original signaling method used on the SCSI-1 bus, uses a common signaling method in which a positive voltage represents a one and a zero voltage (ground) represents a zero, resulting in binary communication. SE signaling is available for any SCSI-1 or SCSI-2 implementation as well as SCSI-3 SPI-1. With most electronic signaling methods, you have a built-in opposition: the faster the transmission speed, the shorter the maximum cable length. Of the three signal types, SE is most susceptible to this limitation. On a Fast-20 bus, for example, an SE cable can only be 5 feet (1.5 meters) long.

HVD

High Voltage Differential (HVD) signaling is also available for any SCSI implementation up to SCSI-3 SPI-1. If you can use SE for the same SCSI implementations, then why use HVD? HVD, as the name implies, uses comparatively more electrical power than the other two types of signaling, resulting in greater allowable bus lengths. Whereas the SE cable is limited to 5 feet (1.5 meters) over Fast-20, an HVD cable can reach 82 feet (25 meters). The "differential" in HVD represents a kind of signaling in which the signal is comprised of the difference in a pair. A one is represented when one wire in the pair transmits a positive voltage, and the other wire transmits zero voltage. The receiving device detects the "difference" between the voltages, and translates it as a one. To transmit a zero, both wires carry a zero voltage. This signaling method is much less susceptible to interference, signal degradation, signal bounce, and crosstalk, allowing HVD to transmit over such a great distance.

It is rare to find HVD in PC servers. It is more often found in older minicomputers and never gained the popularity of SE, primarily because of a higher cost factor.

Caution

Do not mix HVD and SE devices on the same bus. Because of the significantly higher voltage of HVD, SE devices could get smoked—literally.

LVD

Low Voltage Differential (LVD) signaling is the signaling method you are most likely to find today, and it is similar to HVD except, as the name implies, it uses a lower voltage. Advantages of LVD are that you can use both LVD and SE devices on the same bus without electrical hazard, and it allows a longer maximum cable length—up to 39 feet (12 meters). Many SCSI devices are multimode devices; that is, they can operate as either LVD or SE depending on the signaling method of the other devices on the chain. Multimode devices can be indicated in several ways, usually depending on the preference of the company marketing the product. However, it is usually abbreviated as LVD/MSE (where M is multimode) or LVD/SE.

Note the following about using LVD:

- A single SE device among LVD devices on the same chain will drop all devices to SE compatibility. This could have a significant impact if the cable is over 5 feet (1.5 meters) when you add the SE device, because the cable would now be too long.
- Some multimode devices have a jumper setting that forces SE operation. If you want to use LVD, be sure that no devices have enabled this setting.
- Use either LVD or multimode LVD/SE terminators.
- You cannot operate both HVD and LVD on the same bus for reasons of electrical incompatibility, which could damage LVD devices.
- LVD is the only available signaling method for Ultra3, Ultra160, Ultra160+, Ultra320, and probably any other standards that emerge in the near future.

Because the connectors are identical regardless of the signaling method, look for a special symbol on the connector. This will help ensure that you do not mix signaling standards, especially by adding HVD among SE or LVD devices. The signaling symbols appear in Figure 5-10.

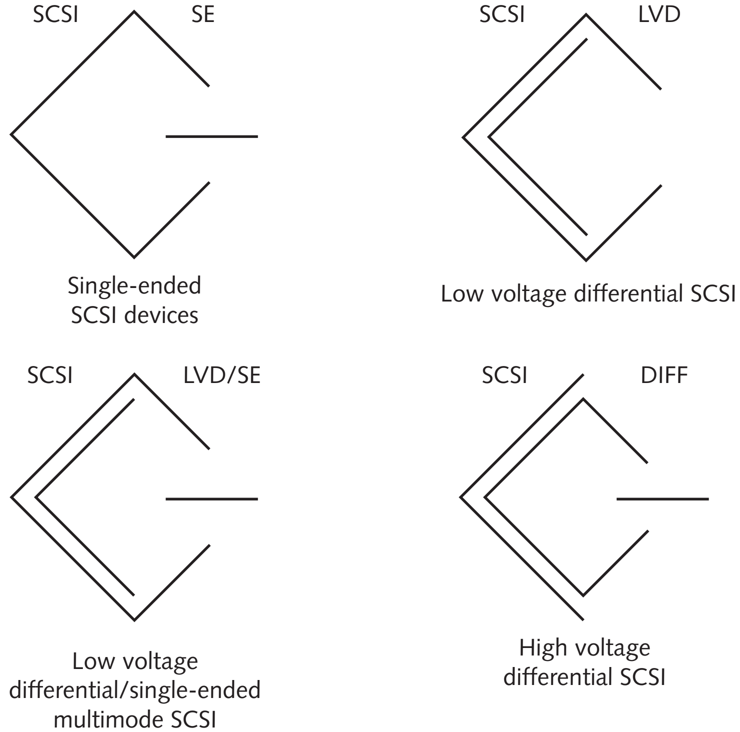

Figure 5-10 SCSI signaling symbols

SCSI Termination

Termination (using terminators at the ends of a SCSI chain) is critical to assure error-free operation on the SCSI chain. A **terminator** absorbs the transmission signal to avoid signal bounce, making it appear to the devices that the cable is of infinite length. Terminators also regulate the electrical load, and are therefore critical in establishing a reliable communications medium. Proper termination requires a terminator at both ends of the SCSI cable. Some devices (and most high-performance host adapters) either automatically terminate or have a setting that allows you to specify termination. Otherwise, obtain a terminator to place over the last connector on the chain. If the last position is in use by a device that does not terminate itself, you can place a terminator over the connection, which allows signal transfer to and from the device while also providing the necessary termination. This is known as **pass-through termination** (see Figure 5-11).

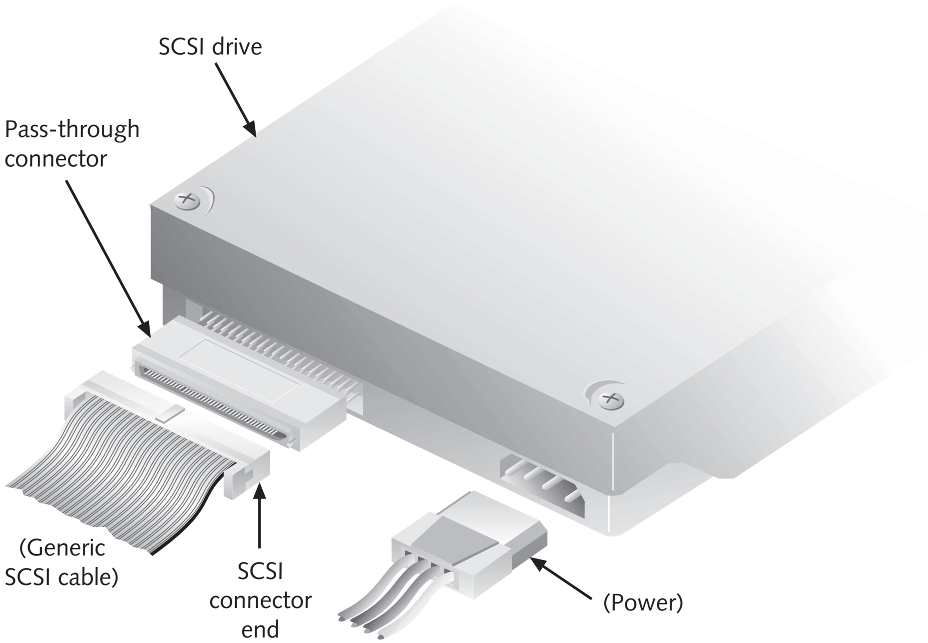

Figure 5-11 Pass-through termination connects the chain on one side of the terminator and the device on the other

Note: Differential drives do not normally include the ability to terminate, so you will have to provide a terminator. You can also purchase cables that have terminators built into the end.

The terminators must also have terminator power supplied, or they cannot properly terminate. If the host adapter does not supply power to the terminator, you can configure a setting on one of the SCSI devices (usually via a jumper) to supply power to a terminator.

TIP: Some administrators configure all SCSI devices to supply power to terminators; that way it never slips through the cracks.

The following list describes the various types of SCSI terminators:

- **Passive termination**—This is the simplest type of termination, but is also the least reliable. Passive terminators use resistors to terminate the SCSI chain, similar to terminators on coaxial Ethernet networks. Passive terminators usually work best on short, SE SCSI-1 buses. It is unlikely you will find many passive terminators in servers.

- **Active termination**—A requirement for faster, single-ended SCSI, active termination adds voltage regulators to provide a more reliable and consistent termination. Another type of active termination is active negation termination, which uses a more complex circuit to stabilize the voltage supply level, further eliminating electrical noise from the signal. Negation terminators are usually only a little more expensive than plain active terminators ($5–10) so I recommend them.
- **Forced perfect termination (FPT)**—In this technology, the termination is forced to a more exact voltage by means of diode clamps added to the terminator circuitry. This advanced form of active termination is very clean, and is the best termination available for an SE bus.
- **HVD termination**—A high voltage bus requires a high voltage terminator. That's pretty much it.
- **LVD termination**—A low voltage terminator for the LVD bus. Many LVD terminators are LVD/SE terminators to accommodate buses with multimode devices. When operating in SE mode, the terminator functions as an active terminator.

Be careful if your server environment mixes narrow SCSI technologies with newer, wider SCSI technologies. First, to place both narrow and wide devices on the bus, you will have to provide an adapter to convert from one width to the other. Converting wide devices to a narrow bus wastes 8 bits of transmission and severely limits the potential of the wide device. On the other hand, converting from 8 bits to 16 bits does nothing to improve the performance of attached 8-bit devices because 8 bits is the highest level of operation anyway. Along with the conversion, you must also consider how termination will affect the two widths. Do you terminate for the 8-bit devices? If so, what about the 16-bit devices? Terminating only for the 8-bits would leave the 16-bit devices with 8 unterminated bits. These dangling bits are called the **high byte** (also known as high 9), and when you have a mix of wide and narrow devices, you can obtain special multimode terminators.

> **TIP**
>
> Multimode terminators flash colors indicating which signaling method is in use. For example, if it's running SE, it might blink yellow, and if it's running LVD, it might blink green. This can be useful in troubleshooting the SCSI chain if you wonder why it seems to be running slow or unreliably over 3-meter lengths—it could be running SE, (just check the light to verify) when you want LVD so that you can benefit from faster speed and longer cable.

Drive Configuration (SCSI ID and LUN)

Knowing the various SCSI technologies, cables, connectors, and termination is fine, but of course you have to also plug it all in! The problem is that with SCSI, misconfiguration can cause many hours of baffling troubleshooting. This section shows you how to correctly configure the devices on the SCSI bus.

Topology

Topology, the physical and/or logical layout of equipment, is simple for SCSI: Use a bus topology because that's all that's available (logically similar to a network bus topology). **Daisy chain** your SCSI devices so that they appear one after the other along the cable. In complicated configurations, you might accidentally connect the bus to itself somewhere along the line, creating a loop. Don't do that or the bus won't work! When configuring a chain from scratch, one end usually connects to the terminated host adapter (unless it's between an internal and external chain) and the other end is terminated as well. The order of the drives (and other SCSI devices, if any) doesn't matter in terms of performance or priority. That's one of the advantages of SCSI—a slower device on the bus doesn't bottleneck all the other devices. In a simple configuration using two SCSI hard disks and a SCSI scanner (see Figure 5-12), the layout is a bus and both ends are terminated.

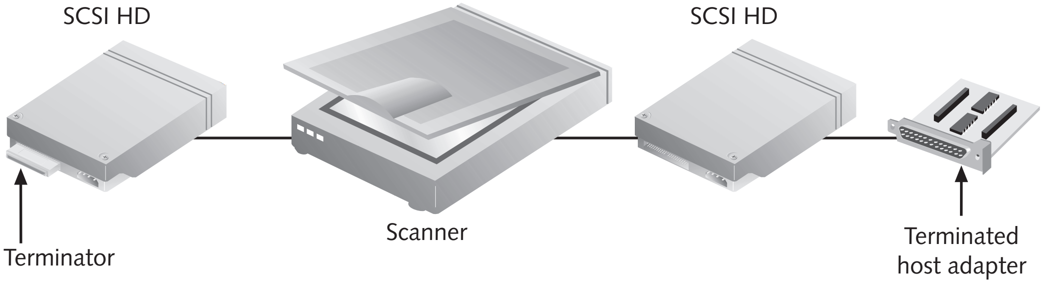

Figure 5-12 A simple SCSI bus topology with terminators at both ends

Realize that although termination takes place at the end of the SCSI bus, the SCSI host adapter is not always at one end of the chain. Because most host adapters have connections for both internal and external devices, a host adapter so connected would appear in the middle of the SCSI chain. In this case, you will have to be sure not to terminate the host adapter (see Figure 5-13).

You are likely to find a SCSI host adapter with multiple channels in many implementations. The primary benefits to multiple channels are twofold: First, having two separate channels allows for twice as many devices. Second, you can configure the two channels to support different signaling technologies. For example, if you have a few SE devices and a few LVD devices, you don't want to mix them or else you'll suffer the performance and length limitations of SE for all the devices. Instead, make one channel SE and the other LVD. That way, the LVD devices can operate at full capacity.

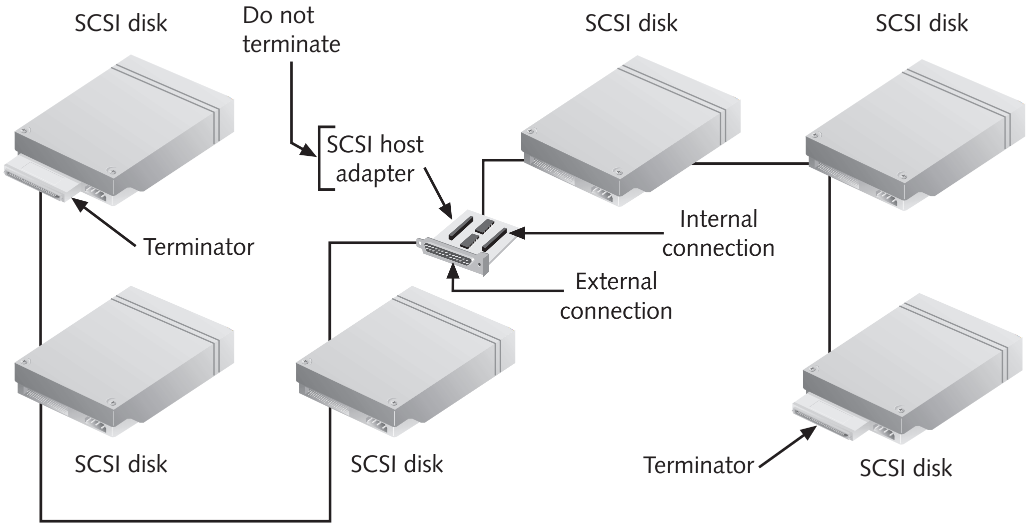

Figure 5-13 Do not terminate a host adapter appearing in the middle of the SCSI chain

> TIP
>
> Reduce inrush power required of the power supply when powering on the server by specifying a power-on delay for the internal SCSI drives. This can be performed via the SCSI BIOS. Similarly, you can delay the powering on of IDE drives in some system BIOS settings. External SCSI drives in their own cabinet enclosure usually have an independent power supply in the cabinet, so power-on delay is not an issue.

SCSI ID Assignment

Each device on the SCSI chain must have a unique SCSI ID number. You configure the device SCSI ID using jumpers, a wheel, or a button on the device, and the adapter is usually preset at ID 7, though you can change it (possibly with a jumper but usually through a software utility). Although you can change the host adapter's SCSI ID, it is highly inadvisable. ID 7 is the highest priority of all the SCSI numbers. The range of available IDs depends on whether you are running narrow (8 IDs) or wide (16 IDs). SCSI ID assignment is important not only to uniquely identify each device, but also to establish which devices have priority when arbitrating over the bus. Arbitration occurs when it must be determined which of two or more devices have control over the bus. Usually, you want to assign a higher priority to slower devices such as scanners or tape drives to make sure that faster devices do not dominate the bus.

The priority of SCSI IDs range from the highest to lowest as follows: 7, 6, 5, 4, 3, 2, 1, 0. On a wide bus, the range is 7, 6, 5, 4, 3, 2, 1, 0, 15, 14, 13, 12, 11, 10, 9, 8. When you install an internal hard disk—simply a matter of screwing it into an available drive bay and plugging in the power and SCSI cables—you also need to configure a SCSI ID using jumper settings. External SCA devices, however, usually configure their own SCSI IDs automatically as soon as you swap the drive into its bay. This is needed to reduce the configuration time, as time is usually of the essence in swapping these drives.

Note

You can divide a single SCSI bus into multiple SCSI **segments** so each device must have a unique ID. However, each segment is electrically independent, and therefore capable of the maximum cable length as if it were truly its own bus. Each SCSI segment requires its own termination.

Along with the unique SCSI ID number is the **logical unit number (LUN)**. The LUN is a subunit of the device, and is used to identify items within the device. For example, a multi-disc CD-ROM changer probably assigns a unique LUN to each disk. If the CD-ROM drive is SCSI ID 5, the first LUN is probably LUN 0, the second is LUN 1, and so forth. Some host adapters do not support LUNs—the absence of LUN support makes the bus scan process during startup faster; so if you need LUN support, make sure the SCSI host adapter supports it. Administrators are not normally concerned with changing LUNs because the device manufacturer defines them.

SCSI Pros and Cons

SCSI has many pros and cons, but as mentioned earlier, the pros far outweigh the cons in most enterprises, and particularly for servers.

Pros

The benefits of SCSI are as follows:

- *Performance*—Aggregating the performance of multiple disks as in a RAID array (covered later in the chapter) generates an appreciable performance gain. The more disks you add, the better the performance (up to the available throughput of the SCSI channel).
- *Expandability*—Several meters of cable length is enough for most SCSI implementations. Most use LVD, which allows for 12 meters. Also, compared to ATA standards that limit you to two disks, the expandability of using up to 15 devices on a SCSI chain is impressive.
- *Redundancy*—Several RAID implementations including RAID-1 and RAID-5 protect your data even if a hard disk fails.

Cons

The drawbacks of SCSI are as follows:

- *Difficult to configure and troubleshoot*—SCSI has so many varying standards, and there are so many opportunities for incompatibility and wrong termination, that it might take a while to get SCSI rolling, especially if you are trying to piece together a SCSI implementation out of existing equipment. If you build from scratch using consistent standards, it's much easier.
- *Expensive*—Smaller organizations or those that have a tight budget might not be able to afford SCSI.
- *Performance*—The performance improvement is not that noticeable if you're only using a single SCSI disk. SCSI shines when using multiple disks for performance and fault tolerance. Otherwise, you're better off saving the money and using ATA-4 or better.

Server administrators usually prefer SCSI because of its advantages in an enterprise environment.

SCSI Summary

In addition to the information in Table 5-3, remember the following about SCSI:

- 8-bit SCSI can connect up to seven devices (actually it's eight devices, but the host adapter counts as one).
- 16-bit SCSI can connect up to 15 devices (actually 16, but the host adapter counts as one).
- Exception to the number of devices: SPI (SCSI-3) Fast-20/Wide can use 15 devices on an HVD cable and seven devices on an SE cable.
- "A" cable is always used for 8-bit SCSI.
- "P" cable is always used for 16-bit SCSI.
- SPI-1 SCSI can use SE at 10 feet (3 meter) lengths, but only 5 feet (1.5 meters) if four or more devices are on the chain.
- Maximum length for all HVD signal cable is 82 feet (25 meters).
- Maximum length for all LVD signal cable is 39 feet (12 meters), but if the chain consists of the host adapter and only one device, you can extend length to 82 feet (25 meters).

Note: SCSI and RAID are troubleshooting studies of their own, and heavily emphasized on the Server+ exam. Properly configuring your SCSI and RAID devices as instructed in this chapter is an important start; however, to account for the many other troubleshooting scenarios you might encounter, refer to a special section, SCSI and RAID Troubleshooting, in Chapter 12.

FIBRE CHANNEL

Fibre Channel (FC) is a storage technology that can use gigabit Ethernet networks, but is primarily intended for fiber optic cable as the name implies. (Fibre Channel technology has roots in Europe, hence the spelling of "Fibre" instead of "Fiber.") FC is a form of storage categorized under the **storage area network (SAN)** umbrella. SAN is a general term that refers to any network-based storage solution that is not server-based. SAN is becoming huge in the server world. Comprehensive coverage of SAN is beyond the scope of this book because it encompasses more than just server issues. In this section, you learn the basics of FC.

First, let's separate out from the discussion another growing storage solution that is sometimes confused with FC—**network attached storage (NAS)**. NAS is one or more storage devices attached to a network, most commonly Ethernet. A NAS device is easy to configure and use, because you just attach it to the rack, plug it into the network, flip its power switch, and voilà! It's ready to use. NAS devices are as accessible as any other device on the network. For smaller organizations, you can use a freestanding NAS device, or for the enterprise, rack models are available. Regardless, when you attach and power up the device, it receives an IP address and identifies the type of network in use. If you want to configure a static IP address, configure security settings or access privileges, or configure RAID, you can launch a management utility (usually browser-based so you can launch it from anywhere in the enterprise) and make changes. NAS devices are very inexpensive compared to FC. NAS devices as well as SAN storage have a significant benefit over traditional file servers in that, well, you don't *have* a file server! This removes an expensive piece of hardware and greatly reduces administration.

In a SCSI FC implementation, a special SCSI host adapter designed for FC connects to the FC bus, which can be up to 10 kilometers (more than 6 miles) in length! That's the benefit of using fiber instead of copper.

FC is run in its own storage "network" of either **Fibre Channel Arbitrated Loop (FC-AL)** or switched fabric. In FC-AL, you can connect 126 storage devices to a special fiber hub in a physical star, logical loop topology. However, all devices on the FC-AL share the available bandwidth, which is OK in many implementations if it meets the needs of the enterprise (see Figure 5-14).

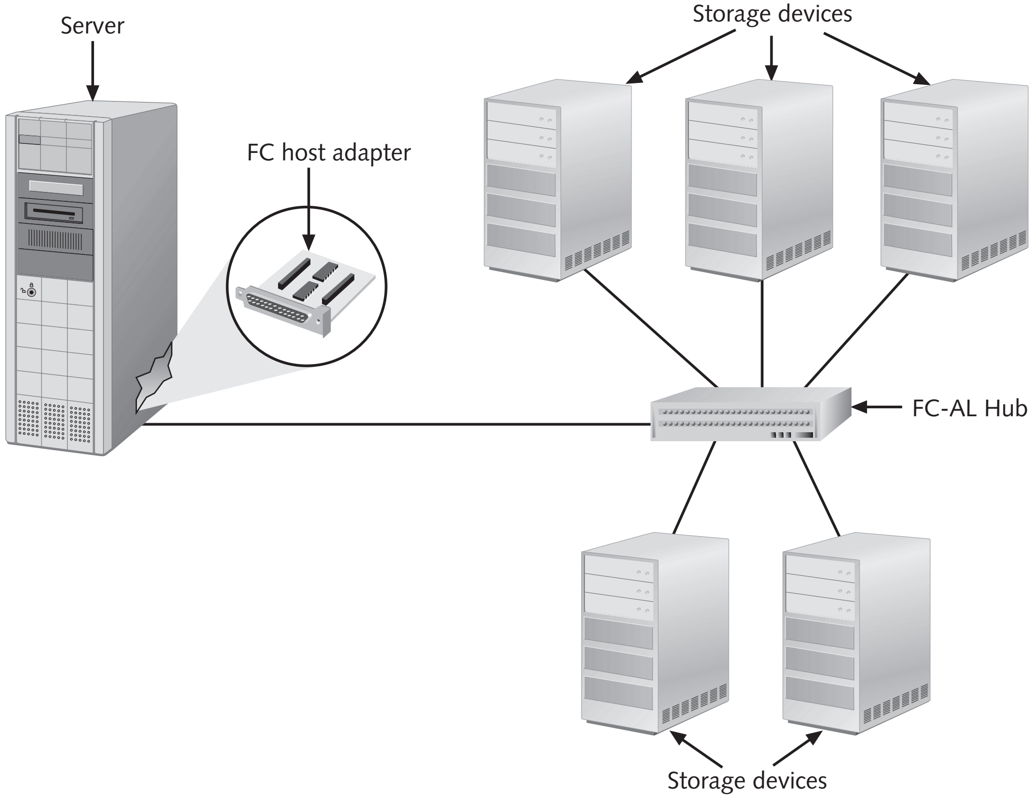

Figure 5-14 An FC-AL configuration connects to a fiber hub and shares available bandwidth

In high-level implementations, you can use a switched fabric to connect to up to 15.5 million storage devices, and the full bandwidth of the channel is available to each device! As you might guess, this implementation is extremely expensive. The term **switched fabric** is somewhat vague in terms of its physical configuration. Simply put, it means that servers storing information on an FC device can reach the device using any number of physical paths. This is analogous to a phone call to a friend in another state. Call her today, and the phone company will route your call using the best path it deems practical. Call her tomorrow, and the phone company might or might not use the same path. However, it doesn't matter to you or to her, so long as the call connects.

Some people have prematurely dismissed FC as too slow at its rated throughput of 100 MBps, especially compared to Ultra160, which can reach up to 160 MBps. In reality, FC is much faster because SCSI cannot sustain speeds of 160 MBps; that's just its burst speed when feeding data from cache. Instead, it will usually dish data out at a sustained rate of around 50 MBps. FC, on the other hand, can sustain data transfer at about 97.5 MBps. The storage media can be whatever file format suits you, and you can use RAID if you like.

Note: It is common to also attach high-speed, high-capacity tape storage to FC.

Typically administrators do not purchase the FC equipment and configure it themselves. Instead, you contract with a SAN storage company specializing in FC to supply the equipment and connect it all together. Maintenance is not usually an issue either, as the storage company also provides maintenance. Larger enterprises often have full-time employees from the storage company on site at all times to maintain and administer the storage and tapes.

The expense of FC is very high, and thus is usually seen in enterprises that need to store terabytes of data, such as in a large data center. The storage devices for a single terabyte of data alone can cost up to $750,000 before installation and costs of other associated equipment, with many implementations totaling millions of dollars.

Note: When receiving FC equipment, be prepared for units weighing upward of a ton and measuring about 7 feet (2.13 meters) high and 10 feet (3.05 meters) wide.

RAID

RAID stands for either Redundant Array of Inexpensive Disks or Redundant Array of Independent Disks. It doesn't really matter—it's just that the history of RAID has changed over the years and the acronym has changed with it. What matters is how you use RAID in the enterprise, and that you understand the characteristics of several levels of RAID, particularly RAID-0, RAID-1, RAID-5, and RAID-0+1.

The main purpose of RAID is to use multiple disks to improve performance, provide redundancy, or both. You configure a RAID array via host adapter hardware or software, though in most cases hardware RAID is preferred for reasons we will discuss at the end of this section. To the operating system and applications, the drives are logically a single drive.

With a controller that supports it, you can obtain a significant performance benefit using a **RAID cache**. Regardless of the version of RAID you use, the host adapter can have a certain amount of memory, usually a minimum of 32 MB, though better adapters have at least 128 MB. The memory is often supported by a battery backup for data integrity, similar to CMOS. The RAID cache fills with data sequentially beyond the actual requested data in anticipation that the next data will soon be requested. If the data is indeed required, the RAID cache serves data more quickly than if data must be retrieved directly from disk.

RAID-0

RAID-0, also known as **disk striping**, lays down data across two or more physical drives. RAID-0 provides no redundancy—if one of the drives fails, all the data is lost, so you must provide another solution for redundancy such as tape backup or RAID-0+1

(see later section). RAID-0 yields a performance advantage because you can aggregate the performance of multiple drives. The SCSI bus can deliver data in a matter of microseconds, but the data transfer of a single disk takes milliseconds, hence a bottleneck at even the fastest drives. RAID-0 helps to mitigate this bottleneck because multiple disks can simultaneously deliver data at once. For example, a RAID-0 array consisting of four drives provides roughly four times the performance of a single drive serving the same data because all four drives on a SCSI chain can operate at once. In Figure 5-15, the letters on each of the drives represent a portion of data, illustrating how the data appears on the drives. In any of the RAID levels, you can further improve performance by dedicating a single drive to each host adapter. (Using two host adapters and one drive on each is known as **duplexing**.) It is possible to use RAID-0 in an IDE implementation, but because IDE can operate only one drive at a time, there is no performance advantage.

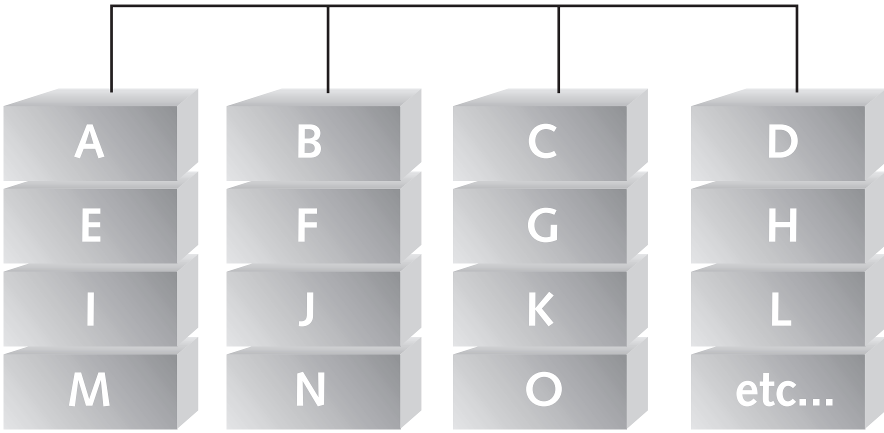

Figure 5-15 This RAID-0 array uses four drives to quadruple hard disk performance

Part of the POST process involves searching for a configurable BIOS from other hardware, particularly SCSI host adapters. To configure a RAID-0 drive once the proper SCSI connections have been made, enter the SCSI adapter's BIOS. Typically, a message appears instructing you to press a keyboard combination to access the BIOS. Then, use the menu system to specify the type of RAID you want to use and save changes.

RAID-1

RAID-1, also known as **disk mirroring**, requires at least two disks to provide redundancy. It can also provide improved performance, but only if using duplexing. In a RAID-1 array, the controller writes the same exact data to two disks at the same time. You can configure multiple adapters, each having two channels with a disk on each channel and capable of simultaneous I/O, to further increase performance (see Figure 5-16).

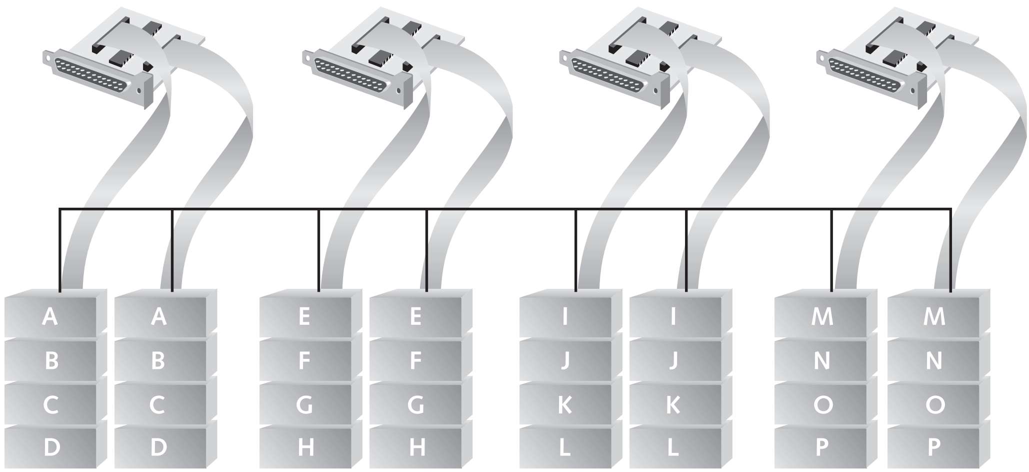

Figure 5-16 RAID-1 using multiple controllers provides both mirroring and high performance

RAID-1 has a higher cost factor or "overhead." If you have two mirrored 40 GB drives, the total raw storage available is 80 GB, but you can only use 40 GB because 40 GB must be available for the mirror. The overhead is 50% because you only use half of the actual disk space. Also, if the drives are not the same size, the mirror is the size of the smaller drive. You can configure RAID-1 through software (using Windows NT/2000, for example), but it requires processor utilization; so in performance terms you're better off using hardware RAID through the controller.

In the event of a failed drive in a RAID-1 array, simply pull out the failed drive and insert a new one. Depending on the method used to configure RAID, you must manually regenerate data from the remaining drive to the new drive through software or BIOS, or the mirror regenerates automatically. With automatic regeneration, the user experience is not disrupted, although performance may suffer temporarily until all data is regenerated to the new drive.

RAID-0+1

As the name implies, **RAID-0+1** offers the best of both worlds: the performance of RAID-0 and the redundancy of RAID-1. In this implementation, two channels and at least four drives are required. Data is striped across two or more disks in the first channel (RAID-0), and the data from the first channel is mirrored to disks in the second channel (RAID-1) in the same striped layout as shown in Figure 5-17. This implementation has a 50% overhead.

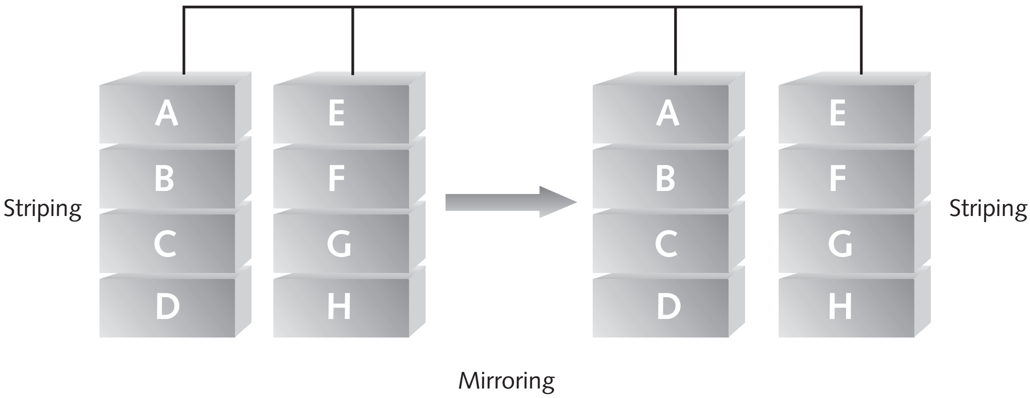

Figure 5-17 A RAID-0+1 array

Note

Contrary to some sources, this is *not* the same as RAID-10. Although RAID-10 also uses both mirroring and striping, one channel mirrors the data and the other channel stripes the same data (see Figure 5-18). In RAID-10, overhead is higher, although two disks could fail and you would still have enough fault tolerance to rebuild the data.

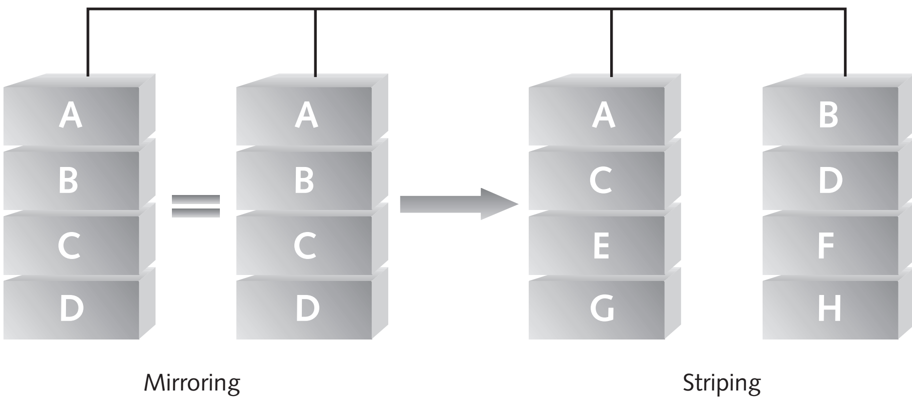

Figure 5-18 A RAID-10 array

RAID-5

RAID-5 offers the performance benefit of RAID-1 striping but also adds redundancy by use of parity. In the case of RAID-5, **parity** is an encoding scheme that represents data appearing on other drives (sometimes called striping with parity). RAID-5 requires at least three hard disks to implement. The host adapter writes data to drives 1 and 2, and on disk 3 writes parity data. At the next write, data writes to drives 2 and 3, and on disk 1 writes parity data, and so on in round-robin fashion (see Figure 5-19). If the first

disk fails, replace it. Then, the system regenerates the data that was supposed to be on the first drive using the parity on disk 2 and disk 3. SCSI is the best choice for this RAID implementation because it offers simultaneous I/O; that is, data reads occur from multiple drives at the same time.

RAID-5 has a lighter overhead in terms of disk space than RAID-1. Overhead is 1/X where X is the number of disks in the array. For example, a five-disk RAID-5 array would have an overhead of 1/5, or 20%. The overhead space is utilized by the parity data.

RAID-5 arrays can take quite a while to rebuild data to a new replacement drive. Also, parity calculation requires CPU cycles, and can significantly affect server performance, especially while rebuilding data to a replacement drive.

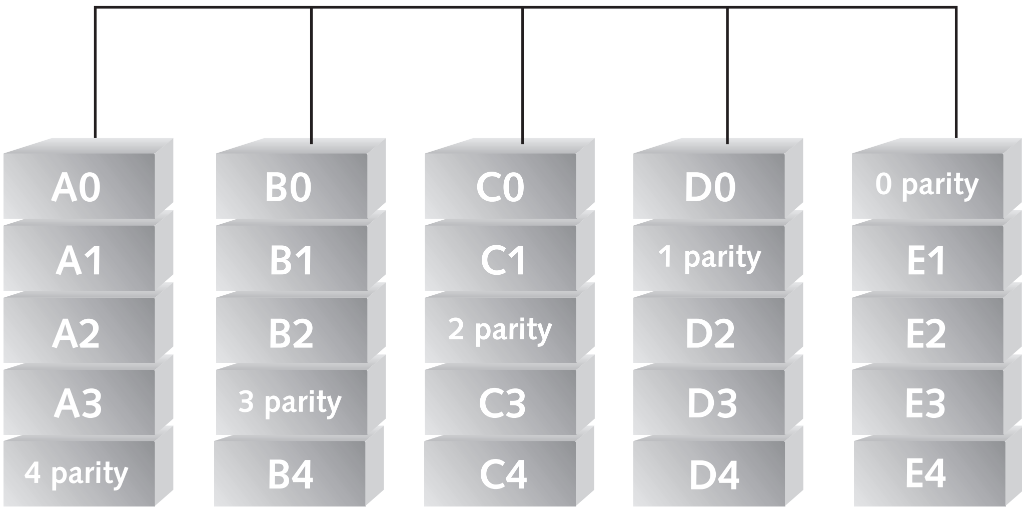

Figure 5-19 RAID-5 writes parity bits across members of the array, and uses the bits to reconstruct data if a drive must be replaced

RAID-1 and RAID-5 reconstruct data with the least amount of disruption if the hard disks are hot-swappable and data is reconstructed on the fly as opposed to rebooting and accessing the BIOS to reconstruct the data.

> **Note** Although manufacturers such as Promise Technology offer EIDE host adapters providing RAID-0, 1, 5, and 0+1, performance benefits are best with SCSI because EIDE cannot perform simultaneous I/O.

Regardless of the RAID configuration you choose (if any), be sure to optimize the hard disk so that it performs at its best. One of the most useful utilities for this purpose is defragmentation software, which arranges the data on your hard disk in a sequential fashion so that reads require less back-and-forth action to locate all the data. (See Chapter 11 for more about defragmentation.)

Software RAID vs. Hardware RAID

Network operating systems such as Windows NT/2000 and Novell NetWare offer software RAID. Software RAID has advantages and disadvantages, as seen in Table 5-4.

Table 5-4 Advantages and Disadvantages of Software RAID

Advantages	Disadvantages
Software RAID is built into the NOS, and is less expensive than hardware RAID.	Less robust than hardware RAID. If the operating system becomes corrupt, the array is at risk.
Generally an easy-to-use graphical interface. You can create a RAID array in a matter of minutes, perhaps without rebooting the system.	Limited configuration capabilities. You are usually limited in the types of RAID configurations you can use. In Windows 2000, for example, you can create a RAID-0, RAID-1, or RAID-5 array but not other RAID configurations such as 10 or 0+1.
	Uses the CPU to perform parity calculations.

Hardware RAID offers advantages and disadvantages, as seen in Table 5-5.

Table 5-5 Advantages and Disadvantages of Hardware RAID

Advantages	Disadvantages
More reliable and robust	More expensive
Faster performance because it is dedicated to RAID functions. Also, parity calculations are handled by hardware RAID (not the main CPU), which is faster than software RAID.	If you need to access the SCSI BIOS, a reboot is required.
The array remains regardless of corruption to the operating system.	

All things considered, experienced administrators usually recommend hardware RAID as a better solution because of its dependability and performance.

Chapter Summary

- Most hard disks include the following major physical components: disk platters, drive heads, head actuator mechanism, spindle motor, logic board, connectors, and jumpers.
- Most current hard disks use between 4 and 11 glass-ceramic platters coated with magnetic media using a thin-film process.
- Drive heads read and write data on the top and bottom of the platter, and glide on a cushion of air when the drive is in operation. When inactive, the heads automatically park.

- An actuator mechanism physically positions the drive heads over locations on the platters determined by the servo mechanism.
- Hard disks can, to a degree, sustain hostile conditions such as shock up to 300 Gs, and last nearly 137 years.
- A primary partition boots to the operating system, while an extended partition exists for the purpose of storing logical drives. A logical drive is a section on the hard disk that appears to the operating system as if it were a separate hard disk and has its own drive letter.
- FDISK is the primary tool for creating and deleting partitions. You can also use third-party tools such as GDISK (from Symantec) or PartitionMagic (from PowerQuest).
- The File Allocation Table (FAT) file system is a Microsoft-based file system compatible with nearly any operating system. FAT has a limit of 2 GB partition size and offers no security. FAT32 offers a 2 TB partition size.
- NTFS, the native file system for Windows NT and Windows 2000, can support extremely large volumes, but it reaches a practical limit at 2 TB. NTFS supports other features such as encryption, compression, and disk quotas.
- HPFS is the operating system of OS/2, and usually coexists with a FAT partition.
- Linux and UNIX can use UFS, NFS, or AFS file systems.
- NetWare uses a NetWare volume in its file system. A NetWare volume is created during installation, and you can add more as necessary. Recent versions of NetWare also offer the NSS file system, which increases performance and storage capacity.
- IDE and ATA are interchangeable terms, though IDE (and EIDE) is actually a marketing term and not a true standard. ATA/IDE operates at a current maximum of 100 MBps and uses no more than two drives per channel. ATA does not support simultaneous I/O.
- ATA cable is either 40-conductor/40-pin or 80-conductor/40-pin. The latter is used for current ATA-4 and ATA-5 implementations.
- Drives can be configured with jumpers to be master or slave. One is not better than the other, except that the master receives preference when the system looks for a boot drive. Cable select makes this easy by automatically detecting that the middle connector is the slave and the end connector is the master.
- ATA drives are inexpensive and easy to configure, but you can only connect two drives to an 18-inch (457.2 mm) cable and the performance is not as good as SCSI.
- What most people call the SCSI "controller" is actually the SCSI host adapter—similar to the inaccurate reference to an IDE "controller."
- SCSI has several versions, each with different benefits and capabilities. Currently, SCSI-3 Ultra320 is the fastest available standard. SCSI uses 8- or 16-bit implementations, A or P cables, and lengths from 5 feet (1.5 meters) to 82 feet (25 meters).

- SCA-2 is the SCSI connection that allows you to plug a hot-swappable drive into an array.
- The SCSI chain is a bus topology and must be properly terminated at both ends using passive, active, forced perfect, HVD, or LVD terminators.
- Never mix HVD devices with LVD or SE devices, as they can damage other equipment.
- SCSI devices must each have a unique SCSI ID. The host adapter is usually ID 7. Devices can also have a subset of identification numbers known as the LUN.
- The priority of SCSI IDs range from the highest to lowest as follows: 7, 6, 5, 4, 3, 2, 1, 0. On a wide bus, the priority is 7, 6, 5, 4, 3, 2, 1, 0, 15, 14, 13, 12, 11, 10, 9, 8.
- Table 5-3 summarizes SCSI specifications.
- NAS is inexpensive network storage that is very easy to install, and you access it from the network (it has an IP address).
- Fibre Channel (FC) is intended for very large storage needs, often in the terabyte range. FC is very expensive but has extremely fast sustained throughput and room for lots of devices.
- All RAID implementations are best served over a SCSI bus, not IDE.
- RAID-0 is striping across at least two drives and offers high performance but no redundancy.
- RAID-1 is mirroring data onto two or more drives simultaneously. Overhead is high at 50%.
- RAID-5 is striping with parity, and is both fast and redundant.
- RAID-0+1 is both mirroring and striping.
- Hardware RAID is usually better than software RAID because it is faster and more reliable.

KEY TERMS

"A" cable — 8-bit, 50-conductor SCSI cable.

active termination — A requirement for faster, single-ended SCSI, active termination adds voltage regulators to provide a more reliable and consistent termination. Another type of active termination is active negation termination, which uses a more complex circuit to stabilize the voltage supply level, further eliminating electrical noise from the signal.

actuator mechanism — A mechanism that physically positions the drive heads to the appropriate location on the disk platter to read or write data.

AFS — Refers to Carnegie-Mellon's Andrew File System, a UNIX file system.

arbitration — Determination of which of two or more devices has control over the bus.

ATA (AT attachment) — Drive technology that attaches to the 16-bit AT bus.

ATA-1 — An ATA standard that supports master, slave, or cable-select determination using jumpers and connecting to a 40-pin cable. The transfer rate is 3.3–8.3 MBps.

ATA-2 — An ATA standard supporting large drive support up to 137 GB. Also known as Fast-ATA-2 and Enhanced IDE (EIDE).

ATA-3 — An ATA standard that supports S.M.A.R.T. and transfer rates up to 16.6 MBps.

ATA-4 — An ATA standard that introduced the optional 80-conductor/40-pin cable and transfer rates up to 33 MBps. Also known as Ultra-DMA and Ultra-ATA. In reference to the transfer rate, you might also see UDMA/33 or Ultra-ATA/33.

ATA-5 — An ATA standard requiring the 80-conductor cable, also adding support for the IEEE-1394 (FireWire) specification and a 66 MBps transfer rate. Later implementations achieve 100 MBps and are also known as Ultra-DMA/100.

ATA-6 — The upcoming official 100 MBps ATA standard.

ATAPI (ATA Packet Interface) — A specification that allows other devices besides hard disks to plug into the ATA interface.

blind connector mating — Refers to the fact that you can't see the SCSI hot-plug connection take place inside the chassis.

cable select — The IDE drive's position on the cable indicates whether it is a master or slave.

command queuing — A method that allows the host adapter to send as many as 256 commands to the drive. The drive stores and sorts the commands for optimum efficiency and performance internally before responding to the host adapter.

cyclical redundancy check (CRC) — A calculation used by the sending device based on the data in the packet. The data arrives at the destination target and another calculation is performed using the same "formula." If the calculation in the packet matches the calculation performed by the destination device, the data is complete and considered error free.

daisy chain — Connecting one device after another on a SCSI bus.

disk mirroring — See RAID-1.

disk platter — A rigid disk inside the sealed hard disk enclosure. Magnetic media on the surface of the platter store the actual hard disk data.

disk striping — See RAID-0.

Distributed File System (Dfs) — A Windows NT/2000 service that deploys what appears to be a single directory structure over multiple physical file servers.

domain validation — The determined SCSI transfer rate is tested, and if errors occur, the rate is incrementally reduced and again tested until no errors occur.

double transition (DT) clocking — Transmitting data on both the rising and falling edges of the clock cycle. On a 16-bit, 40 MHz bus, this yields a transfer rate of 160 MBps.

drive head — A magnetically sensitive device that hovers over the hard disk platter and reads or writes data to the hard disk.

duplexing — Two host adapters with one drive on each adapter.

extended partition — A partition that provides the ability to store logical drives.

Fast SCSI — SCSI operating at 10 MHz instead of 5 MHz.

FAT/FAT32 — A Microsoft-based file system. FAT is capable of 2 GB partitions and FAT32 is capable of 2 TB partitions. Neither file system offers local security features.

FDISK — An MS-DOS utility used to create hard disk partitions.

Fibre Channel (FC)— A storage area network (SAN) SCSI technology that can use gigabit Ethernet networks, but is primarily intended for fiber optic cable as the name implies.

Fibre Channel Arbitrated Loop (FC-AL) — A connection of up to 126 devices on a shared bandwidth fiber hub.

File Allocation Table (FAT) — A Microsoft-based file system compatible with nearly any operating system.

Filesystem Hierarchy Standard (FHS) — A UNIX directory structure to which Linux complies.

FireWire (IEEE 1394) — An extremely fast bus allowing up to 63 connected devices and up to 3200 Mbps throughput in the latest version.

forced perfect termination (FPT) — An advanced form of SCSI termination in which termination is forced to a more exact voltage by means of diode clamps added to the terminator circuitry. FPT is very clean, and it's the best termination available for an SE bus.

FORMAT — A command line utility that creates a Microsoft-based FAT file system.

gray code — A binary code that identifies physical locations on the drive. Gray code is written to the drive by the drive manufacturer.

head crash — When the drive head contacts the disk platter during operation. This can result in corrupt data or damaged hard disk media, especially on older drives.

high byte — The dangling bits resulting from terminating only 8 bits on a 16-bit bus (also known as high 9). Use special terminators that will terminate both the 8 and 16 bits.

High Performance File System (HPFS) — The native file system of IBM OS/2.

high voltage differential (HVD) signaling — SCSI signaling circuitry that uses a comparatively high voltage to extend the length of the SCSI chain to as much as 82 feet (25 meters).

host adapter — The more accurate term for what is usually referred to as an IDE or SCSI hard disk controller. The host adapter is the physical interface between the hard disk and the computer bus.

HVD termination — High voltage termination for HVD signaling.

Integrated Drive Electronics (IDE) — Refers to any hard disk with an integrated controller. Closely associated with the ATA standard.

Journaled File System (JFS) — An OS/2 file system that contains its own backup and recovery capability. Using an indexing system and log to corroborate file changes, JFS can interoperate with the operating system to repair corrupt files.

logical drive — A section on the hard disk that appears to the operating system as if it were a separate hard disk, and that has its own drive letter.

logical unit number (LUN) — A subunit of the SCSI device, used to identify items within the device.

low voltage differential (LVD) signaling — Similar to HVD except for use of lower voltage and shorter cable lengths (39 feet, or 12 meters).

LVD termination — Low voltage termination for LVD signaling.

master (drive) — The drive that receives the first drive letter assignment from the operating system and contains a boot record.

network attached storage (NAS) — One or more storage devices attached to a network, most commonly Ethernet. Simple to configure, you plug in the power, connect it to the network, and turn it on.

NetWare File System — Novell's file system that offers large volume support, efficient cluster size, and local security.

Network File System (NFS) — A UNIX file system that makes files accessible over a network.

Novell Storage Service (NSS) — Operates alongside the traditional NetWare file system to support large files, improve performance, and provide flexible storage management.

NT File System (NTFS) — A Microsoft-based file system designed for Windows NT/2000, offering large volumes and local security.

"P" cable — 16-bit, 68-conductor SCSI cable.

packetization — A data transfer method that reduces the overall communication overhead. Previously, data was transferred over the SCSI bus using a series of phases to set up and transfer data. Packetization streamlines this process by combining the process into a packet, reducing overhead.

parallel bus — A SCSI reference meaning that multiple wires on the cable can transmit data at the same time.

parity — In SCSI, an encoding scheme that represents data appearing on other drives.

passive termination — The simplest type of SCSI termination, but also the least reliable. Passive terminators use resistors to terminate the SCSI chain, similar to the way terminators are used on coaxial Ethernet networks. Passive terminators usually work best on short, SE SCSI-1 buses. It is unlikely you will find many passive terminators in servers.

pass-through termination — If the last position on the SCSI chain is in use by a device that does not terminate itself, you can place a terminator over the connection, which allows signal transfer to and from the device while also providing the necessary termination.

primary partition — A bootable partition on which you can install operating system files.

Quick Arbitration and Selection (QAS) — Reduces overhead by reducing the number of times that arbitration must occur and by allowing a device waiting for bus access to do so more quickly.

RAID-0 — Also known as disk striping, a level of RAID that lays down data across two or more physical drives, benefiting from the combined performance of all drives in the array.

RAID-1 — Also known as disk mirroring, a level of RAID in which the controller writes the exact same data to two disks at the same time (redundancy).

RAID-0+1 — A level of RAID that offers the performance of RAID-0 and the redundancy of RAID-1. In this implementation, two channels and at least four drives are required. Data is striped across two or more disks in the first channel (RAID-0), and the data is mirrored to disks in the second channel (RAID-1).

RAID-5 — A level of RAID that offers the performance benefits of RAID-0 striping but also adds redundancy by use of parity with less overhead.

RAID cache — A high-speed memory cache that fills with data sequentially beyond the actual requested data in anticipation that the next data will soon be requested. If the data is indeed required, the RAID cache serves data more quickly than if data must be retrieved directly from disk.

Redundant Array of Inexpensive (or **Independent**) **Disks (RAID)** — Utilization of multiple disks to improve performance, provide redundancy, or both.

SCSI-1 — The original SCSI implementation.

SCSI-2 — A version of SCSI that introduced Fast and Wide data transmission.

SCSI-3 — A compilation of several different documents, SCSI-3 can be mostly equivalent to SCSI-2 in its features unless several of the various SCSI-3 features are applied. At present, SCSI-3 can be as fast as 320 MBps under Ultra320.

SCSI-3 Parallel Interface (SPI) — See SPI.

SCSI ID — Unique numbering for each SCSI device to ensure proper SCSI operation.

SCSI Interlock Protocol (SIP) — The SCSI-3 parallel command set.

segment — In reference to SCSI, dividing a SCSI bus. Each SCSI segment is electrically independent, and therefore capable of the maximum cable length as if it were truly its own bus. Each segment requires its own termination, and each device must still have a unique SCSI ID across all segments.

servo mechanism — Detects precise cylinder locations on the platter using gray code.

single-ended (SE) signaling — The original signaling method used on the SCSI-1 bus, uses a common signaling method in which a positive voltage represents a one and a zero voltage (ground) represents a zero, resulting in binary communication.

slave (drive) — Equal in every way to the master, except that it does not receive the first drive letter assignment nor contain a boot record.

slack — Space wasted when data does not fill a complete allocation of cluster space.

signaling — Transmission of data using electrical impulses or variations. These electrical transmissions represent data that the sender originates and the receiver translates based upon a mutually agreed-upon method.

SMARTDRV.EXE — An MS-DOS-based caching utility that significantly speeds up file reads and writes.

SPI — SCSI-3 parallel interface, defining SCSI-3 standards in SPI-1 through SPI-3 releases. The original SPI release has been renamed SPI-1 for clarity when comparing against other successive SPI versions. SPI-1 is also known as Ultra SCSI or Wide Ultra SCSI.

SPI-2 — Also known as Ultra2 SCSI and Wide Ultra2 SCSI, a SCSI-3 standard that introduced SCA-2 connectors LVD signaling, and Fast-40 40 MBps transfer rate on a narrow (8-bit) channel or 80 MBps on a wide (16-bit) channel.

SPI-3 — Still in draft stage at the time of this writing, a SCSI-3 standard that introduces CRCs for data integrity, domain validation, DT clocking, packetization, and QAS. Also known as Ultra3 SCSI.

SPI-4 — The latest SCSI-3 specification, still in draft form at the time of this writing. Most hard disk and host adapter manufacturers have products using the standard's 320 MBps data rate. This data rate is accomplished by doubling the bus speed from 40 MHz to 80 MHz and using DT clocking. Manufacturers are calling this standard Ultra320.

storage area network (SAN) — Generally refers to Fibre Channel and any other type of network-based storage solution that is not server-based.

switched fabric — A somewhat inexact reference to the connection to the FC storage. The connection can use any number of connection routes, depending on which one is deemed best at that particular moment.

terminator — A connector placed the end of a SCSI chain that absorbs the transmission signal to avoid signal bounce, making it appear to the devices as if the cable was of infinite length. Terminators also regulate the electrical load, and are therefore critical in establishing a reliable communications medium. Proper termination requires a terminator at both ends of the SCSI cable.

thin-film — A magnetic medium applied to disks in a near perfect, continuous vacuum.

traditional NetWare file system — Offers similar, competitive features to the NTFS file system, and allows you to create NetWare volumes.

Ultra160, **Ultra160+** — A collection of SCSI-3 standards that ensure compliance with a minimum level of SCSI-3 standards, offering speeds of 160 MBps.

UNIX File System (UFS) — The UNIX file system, which supports large volumes and local security.

voice coil — A construction used by the hard disk actuator mechanism to move from one location to the next.

volume — In NetWare, a collection of files, directories, subdirectories, and even partitions.

Wide SCSI — Utilizing 16-bit transfer instead of 8-bit (which is "narrow" SCSI).

Review Questions

1. Which of the following are physical hard disk components? (Choose all that apply.)
   a. disk platters
   b. spindle motor

c. compact disc

d. head actuator mechanism

2. A head crash is:

a. the drive head colliding with the sprocket

b. the drive head colliding with the disk platter

c. the drive head in contact with a stationary platter

d. the inevitable conclusion to a wild party

3. A primary partition:

a. is contained in an extended partition

b. contains logical drives

c. always represents the entire hard disk

d. can contain a bootable operating system

4. Which of the following is not a hard disk utility?

a. GDISK

b. FDISK

c. FORMAT

d. TETRIS

5. What is "slack" on a hard disk?

a. corrupt data on the physical media

b. cluster space that is only partially occupied

c. excessive free space

d. unterminated SCSI signals

6. The better technical term for IDE is:

a. EIDE

b. ATA

c. SCSI

d. FC-AL

7. Which two ATA implementations can use 100 MBps?

a. ATA-3

b. ATA-4

c. ATA-5

d. ATA-6

8. What is an advantage of IDE/ATA?
   a. up to 15 devices on a single channel
   b. connects to a very fast switched fabric
   c. inexpensive
   d. up to 82 feet (25 meter) cable length
9. What is a synonym for an IDE or SCSI controller?
   a. host adapter
   b. array control unit
   c. domain controller
   d. SCA connector
10. What does a SCSI parallel bus allow?
    a. arbitrated I/O
    b. simultaneous I/O
    c. LVD
    d. multiple channels on a single controller
11. Which version of SCSI introduced both Fast and Wide implementations?
    a. SCSI-1
    b. SCSI-2
    c. SCSI-3
    d. SCSI Fast and Wide have been formally introduced as a standard.
12. What is command queuing?
    a. a method that allows the host adapter to send up to 256 commands to the drive
    b. ordering of I/O operations on the IDE host adapter
    c. a memory queue that caches sequentially read hard disk data
    d. a management system that reorders unprocessed commands to the host adapter
13. A Wide SCSI implementation would use which of the following items?
    a. 8-bit cable
    b. 50-pin connector
    c. 68-pin connector
    d. 16-bit cable
14. What does an SCA-2 connector allow you to do?
    a. replace disk platters on the fly
    b. hot swap hard disks

c. hot swap terminators

d. convert from a 50-pin device to a 68-pin cable connector

15. How does double transition (DT) clocking improve SCSI performance?

a. transmits data twice per second

b. transmits data on the rising and falling edges of the clock cycle

c. fills the RAID cache during idle periods

d. relieves the main CPU of parity processing burdens

16. Why shouldn't you mix HVD devices with LVD devices?

a. The LVD devices can be damaged by the higher voltage.

b. The HVD devices will block communication on the bus.

c. The HVD devices will shorten the overall cable length.

d. The LVD devices can damage the HVD devices.

17. What is the best form of termination for an SE bus?

a. LVD termination

b. HVD termination

c. passive termination

d. FPT

18. What is the SCSI ID of most host adapters?

a. 0

b. 1

c. 7

d. 15

19. Which of the following is not a benefit of Fibre Channel?

a. simple installation

b. extremely high performance

c. extremely high storage quantity

d. extremely high number of connected devices

20. If you want both high performance and redundancy, which two RAID implementations would be suitable?

a. RAID-0

b. RAID-1

c. RAID-5

d. RAID-0+1

Hands-on Projects

Project 5-1

In this project, you will inspect the physical hard disk components.

Your instructor should have an old scrap hard disk for you to examine. If necessary, unscrew the cover (ask for a special torx screwdriver if necessary). Examine the inside of the disk and answer the following questions:

1. What color is the platter? What does this tell you about how it was manufactured?
2. How many platters are there?
3. Observe the heads. Are they in contact with the platter?
4. Optionally, plug the drive into a power connector with the cover off and watch the action of the drive. Don't touch the components while plugged in and running.

Project 5-2

In this project, you will install two IDE hard disks on a channel.

Connect a hard disk to a computer (it can be a desktop PC or a server) as follows:

1. Take the case off the computer.
2. Locate the IDE connectors. They are probably on the motherboard, but might be on an add-in host adapter. On the motherboard, they are usually labeled in the circuit board as IDE 1 and IDE 2 or IDE 0 and IDE 1.
3. Connect the black end of an 80-conductor, 40-pin IDE cable to the IDE connector. Verify correct orientation using the stripe on the cable, a notch, and a missing pin 20 (refer to Figure 5-3). Depending on the manufacturer, not all these features may be available, but there should be enough to orient the cable correctly. The objective is to line up Pin 1 with the highlighted stripe on the cable (usually red or blue). Pin 1 on the motherboard might be labeled with a number "1" or a triangular arrow.
4. On the drive, configure the jumpers for cable select by placing a jumper over the appropriate pins. Usually, there is a CSEL marking on the drive that indicates which pins you need to jumper. If there is insufficient information, ask your instructor for direction.
5. Optionally, install the drive into a drive bay in the computer. This is not strictly necessary because the drives will be grounded through the power connector. However, do not place the drive anywhere such that the circuit board on the hard drive is in contact with any other metallic components in order to avoid electrical damage. Your instructor will advise you.
6. Connect the IDE cable to the drive. The master drive goes on the opposite end from the host adapter IDE connection, and since there is only one drive, connect it to the master position. Again, verify proper alignment using the stripe, notch, and missing Pin 20. As a rule, the stripe on the cable is adjacent to the power supply connector on the hard drive.

7. Plug in power to the hard drive. The power connector is keyed so it will only fit using the proper orientation.

Project 5-3

In this project, you will access the BIOS configuration to view details about a drive.

1. Turn on the power to the computer. Access the BIOS (look for a message telling you how or ask the instructor).
2. Navigate the menu system and locate information about the drive. Who is the manufacturer? How many megabytes are on the drive? (Not all BIOS will allow you to view this information. If so, proceed to Step 3.)
3. Navigate to where you can adjust the boot order. If necessary, adjust the boot order to boot from the floppy disk first, then the hard disk.
4. Exit the BIOS, saving changes.

Project 5-4

In this project, you will boot to a Windows 98 boot floppy, use FDISK to partition the drive, and FORMAT to format the drive. (If you are familiar with FDISK and FORMAT, you can skip this project.)

1. Place a Windows 98 startup disk into the floppy drive.
2. Turn on the computer, and wait for the boot floppy to display a menu. Select the option to **Start computer without CD-ROM support** and press **Enter**. (If the system starts with CD-ROM support, that's OK; it just takes longer to boot and we don't need the CD-ROM right now.)
3. Type **FDISK** at a command prompt and press **Enter**.
4. You are prompted for whether you want to enable large disk support. Select **Y** and press **Enter**. (If an NTFS partition is already on the disks, also answer **Y** to the next prompt to treat NTFS partitions as large.)
5. An FDISK Options menu appears. Select option **4** to view the partition information (there might not be anything). If there are existing partitions, what types are there?
6. Press the **Esc** key to return to the FDISK Options menu.
7. If there are partitions on the disk, select option **3**, Delete partition or Logical DOS Drive. If there are no partitions, proceed to step 10.
8. You must delete logical drives before deleting an extended DOS partition. If you try to delete a partition prematurely, the FDISK messages will tell you what you need to do. Delete all partitions, deleting the primary DOS partition last.
9. Press the **Esc** key to return to the FDISK Options menu.
10. Enter **1** to create an MS-DOS partition or Logical DOS Drive.
11. Enter **1** again to Create a Primary DOS Partition. FDISK checks the drive integrity before proceeding.

12. A prompt asks if you want to use the maximum available space; enter **N** and press **Enter**. You do not want to use the entire space so that you can create extended partitions and logical drives. Your instructor will guide you as to how much space to use.
13. FDISK verifies drive integrity again. Then enter the amount of space to create and press **Enter**. The partition is created.
14. Press **Esc** to return to the main FDISK Options menu.
15. You must set the partition to active, or it will not be bootable. Select option **2** to do this, and enter a **1** to specify the partition you want to make active.
16. Press **Esc** to return to the main FDISK Options menu, and **Esc** again to exit FDISK.
17. Now you must reboot the computer, or else the FORMAT command in the next step will not work. The Windows 98 startup disk should still be in the floppy drive.
18. At the boot menu, again start without CD-ROM support.
19. Type **FORMAT C: /S** (the /S switch makes the drive a system, or bootable disk).
20. A warning informs you that all data on the drive will be lost. Type **Y** and press **Enter**. The FORMAT process might take a few minutes depending on the size of the partition.
21. After the format is complete, you can enter a volume label of up to 11 characters. Type **Master** and press **Enter**. The drive is now available for use.

Project 5-5

In this project, you will configure a SCSI host adapter and two or three SCSI hard drives.

1. We are switching gears from IDE to SCSI—so with the power off, remove the IDE host adapter (if installed as a card) and IDE hard disk.
2. Read any identifying information from a SCSI host adapter provided by your instructor. Who is the manufacturer? What is the model number? Does it specify its capabilities (bit width, signaling method, etc.)? Optionally, you might also browse to the manufacturer's web site to glean more information about this specific adapter.
3. Install the SCSI host adapter into an available PCI slot. (For specifics about PCI card installation, see Chapter 6.)
4. Plug the SCSI drives into any available connector on the provided SCSI cable but not the last connector. If the cable is round, then it's for external use. If it's flat ribbon, then it's for internal use. Plug the other end of the cable into the host adapter connector. (The equipment provided should all match compatible specifications.) Before proceeding, do you think it is appropriate to turn on power and use the hard drive? Why or why not?
5. In answer to the question above, you should not turn on the system yet because nothing has been done about termination. A terminator is supplied by your instructor. Identify the terminator and write down what type it is: HVD, LVD, SE, mixed, passive, etc.

6. Plug the terminator into the last connector on the cable.
7. Ask your instructor about terminating the host adapter, as the method might vary from one adapter to another. Most likely, you can specify termination in the BIOS settings, which we'll look at in the next project.

Project 5-6

In this project, you will access the SCSI host adapter bus and configure a RAID solution.

1. Confirm correct termination and connections on the SCSI drive and host adapter. Then, power on the system.
2. After the system POST, you should see a BIOS screen for the SCSI host adapter. It might not be visible for long, so stay sharp. Look for instructions that indicate how to access the SCSI BIOS (often for Adaptec controllers, press Ctrl+A). If you can't find any, your instructor will inform you as to how to access the SCSI BIOS.
3. Once in the BIOS, use the menu system to verify that the drive appears correctly on the bus.
4. Using the SCSI BIOS system, configure the RAID array. Experiment and create what you like. The BIOS probably offers several choices. Consider creating a RAID-0 or RAID-5 array if you have three drives, or a RAID-1 if you have only two drives. Building the arrays could take quite a while with larger drives.

Note

Note that in an actual implementation, if the boot device will be part of the array, you must create it prior to installing the operating system.

Case Projects

1. Sal, an ad-hoc administrator from another site, cannot understand why the hard disks on his newly configured IDE controller do not function. He says that the IDE controller is compatible with Ultra-DMA/100, although the two brand-new hard disks he placed on a single channel are both Ultra-DMA/66. Sal is certain that he has the right cable because he counted all 40 conductors and knows that he needs 40 pins to make the connection. When he powers up the computer, a POST message appears stating that he has no operating system. What should you tell Sal?
2. You have again volunteered your services to the nonprofit organization, KidHelp. KidHelp has been growing rapidly and has developed a database that is over a terabyte in size. Currently, KidHelp uses several file servers that collectively store the database, but they want to centralize into a single storage solution. A wealthy donor has told them to procure the best solution regardless of cost. Also, they want fast, regular backups and immediate protection in the event of a hard drive failure. What solution might you suggest to KidHelp?